

Neuropsychology of reinforcement processes in the rat

*A dissertation submitted for the degree of
Doctor of Philosophy*

Rudolf Nicholas Cardinal
St John's College, Cambridge

January 2001

Department of Experimental Psychology
University of Cambridge
Downing Street
Cambridge CB2 3EB
United Kingdom

To Ann & John

'How the purer Spirit is united to this clod, is a knot too hard for fallen Humanity to untie.'

Joseph Glanvill (1661), *The Vanity of Dogmatizing*.

'It ain't necessarily so.'

Ira Gershwin and DuBose Heyward (1935), *Porgy and Bess*.

Contents

<i>Preface</i>	11
<i>Abstract</i>	12
<i>Acknowledgements</i>	14
<i>Publications</i>	15
Abstracts	15
Papers and book chapters	15
<i>Abbreviations used in this thesis</i>	16
CHAPTER 1. INTRODUCTION	19
OVERVIEW	19
PSYCHOLOGICAL MECHANISMS FOR ACTION	19
<i>Unlearned behaviour</i>	19
<i>Pavlovian conditioning</i>	20
Representations formed during Pavlovian conditioning	20
<i>Goal-directed behaviour</i>	23
Instrumental contingency	24
Incentive value	24
Hedonic assessment: an implied, distinct valuation process	24
Discriminative stimuli	25
<i>Stimulus–response habits</i>	25
<i>Pavlovian to instrumental transfer</i>	26
Response-specific potentiation	26
General potentiation of behaviour	26
Simple transfer tasks	28
<i>Summary</i>	28
CONDITIONED REINFORCEMENT	30
<i>Validity of conditioned reinforcement as a concept</i>	30
<i>Relationship of conditioned reinforcement to other processes controlling instrumental performance</i>	31
NEURAL DISSOCIATIONS WITHIN ASSOCIATIVE LEARNING	32
<i>Overview of the limbic corticostriatal circuits considered in this thesis</i>	32
<i>Interpretation of lesion studies</i>	34
<i>A summary of candidate neural structures involved in Pavlovian conditioning</i>	36
Stimulus representation: sensory thalamus, primary and higher-order sensory cortices	36
The amygdala	36
Conditioning of simple CS–UR skeletal responses with high temporal precision: the cerebellum	40
The anterior cingulate cortex and stimulus–reinforcer associations	41
Expression of Pavlovian conditioning: amygdala–accumbens and cingulate–accumbens interactions	42
The anterior cingulate cortex: unanswered questions	43
<i>A summary of candidate neural structures that influence instrumental performance</i>	43
Instrumental performance: outside the ventral striatum	43
The nucleus accumbens, dopamine, and the impact of Pavlovian conditioned stimuli	46
DELAYED REINFORCEMENT	51

<i>Delayed reinforcement in learning</i>	52
<i>Choice, and pathological choice, from the perspective of utility theory</i>	52
Utility theory.....	52
Pathological choice in the context of utility theory.....	53
Views of choice in the brain	54
<i>Impulsivity and impulsive choice</i>	56
<i>Delayed reinforcement in choice</i>	57
<i>Neurochemical and neuroanatomical studies of delayed reinforcement</i>	58
Serotonin (5-HT).....	58
Dopamine and attention-deficit/hyperactivity disorder (ADHD)	59
The prospect of delineating neural circuitry involved in choice of delayed reward.....	59
<i>Tasks used to study choice of delayed reinforcement</i>	60
<i>Conditioned reinforcement in choice of delayed rewards</i>	61
ORGANIZATION OF EXPERIMENTAL WORK IN THIS THESIS	61
CHAPTER 2. GENERAL METHODS	63
SUBJECTS AND HOUSING CONDITIONS	63
SURGERY.....	63
<i>General surgical technique</i>	63
<i>Excitotoxic lesions</i>	63
<i>Implantation of bilateral intracranial cannulae</i>	65
HISTOLOGICAL ASSESSMENT	66
<i>Perfusion and tissue fixation</i>	66
<i>Nissl staining with cresyl violet</i>	66
<i>Immunocytochemical staining for neuronal nuclei</i>	66
DEFINITION OF REINFORCEMENT SCHEDULES	67
BEHAVIOURAL APPARATUS	67
DATA ANALYSIS	68
CHAPTER 3. ROLE OF THE ANTERIOR CINGULATE CORTEX IN THE CONTROL OVER BEHAVIOUR BY PAVLOVIAN CONDITIONED STIMULI	70
INTRODUCTION	70
<i>Delineation and connections of the anterior cingulate cortex (ACC) in the rat</i>	70
<i>Involvement of the rat ACC in stimulus–reinforcer association</i>	73
<i>Lesion methods and sites within the ACC</i>	75
EXPERIMENT 1: EFFECTS OF ACC LESIONS ON TEMPORALLY DISCRIMINATED APPROACH, RESPONDING FOR CONDITIONED REINFORCEMENT, AND FEAR CONDITIONING TO A DISCRETE CUE	76
<i>Methods</i>	76
Overview	76
Temporally discriminated approach	76
Acquisition of a new response with conditioned reinforcement	76
Intracranial infusion during conditioned reinforcement test.....	77
Autoshaping.....	77
Sucrose consumption.....	78
Locomotor activity in a novel environment	78

Fear conditioning to a discrete cue	79
<i>Results</i>	79
Histology	79
Temporally discriminated approach	84
Responding for conditioned reinforcement	85
Autoshaping	86
Sucrose consumption	87
Locomotor activity in a novel environment	88
Freezing to an aversive CS	88
Summary	88
<i>Discussion</i>	89
EXPERIMENT 2: EFFECT OF ACC LESIONS ON AUTOSHAPING PERFORMANCE	92
<i>Methods</i>	92
<i>Results</i>	92
Histology	92
Pre-operative acquisition	94
Post-operative performance	94
Probe test	95
Omission test	96
Locomotor activity in a novel environment	96
Summary	96
<i>Discussion</i>	97
EXPERIMENT 3: EFFECTS OF ACC LESIONS ON ‘SIMPLE’ PAVLOVIAN–INSTRUMENTAL TRANSFER	100
<i>Methods</i>	100
Subjects	100
Simple Pavlovian to instrumental transfer	100
<i>Results</i>	101
Pavlovian training	101
Instrumental training	101
Transfer test	102
<i>Discussion</i>	102
EXPERIMENT 4: EFFECTS OF ACC LESIONS ON A TWO-STIMULUS DISCRIMINATED APPROACH TASK	103
<i>Methods</i>	104
Overview	104
Two-stimulus temporally discriminated approach task	104
Two-stimulus test of conditioned reinforcement	104
<i>Results</i>	104
Histology	104
Two-stimulus discriminated approach task	108
Two-stimulus conditioned reinforcement test	111
Summary	112
<i>Discussion</i>	112
GENERAL DISCUSSION	113
<i>Contribution of the ACC to instrumental and Pavlovian behaviour</i>	113

<i>A time-limited role for the ACC?</i>	113
<i>Synthesis: a suggested role for the ACC in ‘disambiguating’ stimuli for its corticostriatal circuit</i>	115
<i>Interactions of the ACC with the amygdala and perirhinal cortex</i>	117
<i>Comparison with other interventional studies in rodents</i>	118
<i>Homology between rodent and primate ACC</i>	120
<i>Interventional studies in primates</i>	122
<i>Correlational studies in humans</i>	122
Emotional states, emotionally significant stimuli, and mood	122
Attention, conflict monitoring, error detection, and action selection	123
<i>Relating human and rodent studies</i>	127

CHAPTER 4. ROLE OF THE NUCLEUS ACCUMBENS CORE AND SHELL IN PAVLOVIAN– INSTRUMENTAL TRANSFER..... 128

INTRODUCTION	128
METHODS	129
‘Response-specific’ and ‘general’ Pavlovian–instrumental transfer tests	129
RESULTS	131
Histology	131
Pavlovian training	138
Instrumental training	138
Response-specific PIT	139
Retraining	141
General PIT	142
DISCUSSION	144
<i>The psychological basis of response-specific PIT</i>	144
<i>The contribution of the <i>Acb</i> to PIT</i>	145
<i>The relationship between PIT and conditioned reinforcement</i>	146

CHAPTER 5. LOCAL ANALYSIS OF BEHAVIOUR IN THE ADJUSTING-DELAY TASK FOR ASSESSING CHOICE OF DELAYED REINFORCEMENT 148

INTRODUCTION	148
EXPERIMENT	149
<i>Methods</i>	149
Subjects	149
Adjusting-delay technique for assessing choice with delayed reinforcement	149
Analysis of behavioural data	150
<i>Results</i>	152
COMPUTER SIMULATIONS	156
<i>Methods</i>	156
<i>Results</i>	159
Local analysis of the simulated decision rules	159
Achievement of stability criteria by a delay-independent decision rule	163
Effect of bias on dB’ using a delay-independent decision rule	164
DISCUSSION	164
<i>Interpretation of cross-correlational analysis</i>	164

<i>Stability does not imply sensitivity to the adjusting delay</i>	166
<i>Possible reasons for the present failure to observe sensitivity to dB</i>	166
<i>Effects of manipulations that alter subjects' preferences in a delay-independent manner</i>	167
Effects of extrinsic manipulations	167
Effects of dA on dB'	167
<i>Comparison to free-operant schedules of reinforcement</i>	168
SUMMARY	169
CHAPTER 6. THE EFFECTS OF D-AMPHETAMINE, CHLORDIAZEPOXIDE, ALPHA- FLUPENTHIXOL AND BEHAVIOURAL MANIPULATIONS ON CHOICE OF SIGNALLED AND UNSIGNALLED DELAYED REINFORCEMENT	170
INTRODUCTION	170
METHODS	171
Subjects, apparatus, and behavioural task	171
Systematic technique for assessment of preference for delayed reinforcement	172
Pharmacological and behavioural manipulations	174
Statistical analysis	176
RESULTS	176
1. <i>Acquisition and baseline performance</i>	176
Acquisition of sensitivity to delay	176
Effect of cues on speed of acquisition	176
2. <i>Baseline performance</i>	176
Effect of cues on choice (between-subjects comparison).....	176
Omissions and latencies	177
3. <i>Pharmacological manipulations</i>	178
Effects of d-amphetamine.....	179
Effects of chlordiazepoxide	181
Effects of α -flupenthixol.....	183
4. <i>Behavioural manipulations</i>	185
Omission of delays.....	185
Effect of cues on choice (within-subjects comparison).....	186
Effects of prefeeding	187
Descending delays.....	188
Extinction	189
DISCUSSION	190
<i>Task validation</i>	190
<i>Role of signals present during the delay</i>	191
<i>Effects of d-amphetamine</i>	192
<i>Effects of chlordiazepoxide</i>	193
<i>Effects of alpha-flupenthixol</i>	194
<i>Conclusions</i>	194
CHAPTER 7. CONTRIBUTIONS OF LIMBIC AND PREFRONTAL CIRCUITRY TO CHOICE OF DELAYED REINFORCEMENT	195
INTRODUCTION	195

EXPERIMENT 1. EFFECTS OF LESIONS OF THE ANTERIOR CINGULATE CORTEX.....	197
<i>Methods</i>	197
<i>Results</i>	197
Histology.....	197
Pre-operative acquisition.....	200
Baseline post-operative performance.....	200
Effect of omitting all delays.....	200
<i>Summary</i>	201
EXPERIMENT 2. EFFECTS OF LESIONS OF MEDIAL PREFRONTAL CORTEX.....	203
<i>Methods</i>	203
<i>Results</i>	203
Histology.....	203
Body mass.....	206
Pre-operative acquisition.....	206
Baseline post-operative performance.....	206
Effect of omitting all delays.....	207
Locomotor activity in a novel environment.....	207
<i>Summary</i>	207
EXPERIMENT 3. EFFECTS OF LESIONS OF THE NUCLEUS ACCUMBENS CORE.....	209
<i>Methods</i>	209
Food consumption tests.....	209
<i>Results</i>	209
Histology.....	209
Body mass.....	211
Pre-operative acquisition.....	211
Baseline post-operative performance.....	211
Effect of omitting all delays (ABAB design).....	213
Effect of omitting all delays (AAABBB design).....	213
Prolonged training without delays, and subsequent reintroduction of delays.....	214
Locomotor activity in a novel environment.....	216
Food consumption tests.....	216
<i>Summary</i>	217
EXPERIMENT 4. EFFECTS OF INTRA-ACCUMBENS AMPHETAMINE ON CHOICE OF SIGNALLED AND UNSIGNALLED DELAYED REINFORCEMENT.....	218
<i>Methods</i>	218
Abbreviated delayed-reinforcement task for intracranial infusions.....	218
<i>Results</i>	218
Histology.....	218
Acquisition and baseline performance.....	220
Re-establishment of baseline performance following surgery.....	220
Effects of intra-accumbens amphetamine on choice.....	220
Effects of intra-accumbens amphetamine on latencies and nose-poking behaviour.....	222
<i>Summary</i>	223
DISCUSSION.....	224

<i>Effects of anterior cingulate cortex lesions</i>	224
<i>Effects of medial prefrontal cortex lesions</i>	225
Contingency perception.....	225
Timing ability	225
<i>Effects of nucleus accumbens core lesions</i>	225
Primary motivational changes.....	226
Altered sensitivity to reinforcer magnitude or delay?	226
Hyperactivity and impulsivity: behavioural comparison to models of ADHD	229
Implications for theories of nucleus accumbens function	230
<i>Effects of intra-accumbens amphetamine</i>	231
<i>Autoshaping and impulsivity</i>	232
<i>The possible role of other structures connected to the nucleus accumbens core</i>	233
<i>Conclusions</i>	234
CHAPTER 8. GENERAL DISCUSSION	235
<i>Introduction</i>	235
<i>Summary of results</i>	235
Role of the anterior cingulate cortex in Pavlovian conditioning.....	235
Role of the nucleus accumbens core and shell in response-specific Pavlovian-instrumental transfer	235
Behavioural tasks used to assess preference for delayed reinforcement.....	236
Effects of d-amphetamine, α -flupenthixol, and chlordiazepoxide on preference for signalled and unsignalled delayed reinforcement.....	236
Neural basis of preference for delayed reinforcement	236
<i>Anterior cingulate cortex function</i>	237
<i>Theories of learning and choice with delayed reward</i>	239
<i>Theories of nucleus accumbens function and the neural basis of delayed reward</i>	241
1. The striatum as a switching device	241
2. Acute modulation of striatal function by dopamine	244
3. The striatum and learning	244
4. Dopaminergic effects on striatal learning; implications for addiction.....	245
5. Incorporation of the present findings relating to delayed reinforcement	247
<i>Reinforcement learning in the brain: an integrative view</i>	248
REFERENCES	250

Preface

The following work was carried out at the Department of Experimental Psychology, University of Cambridge, during the years of 1997–2000, under the supervision of Professor Barry J. Everitt.

I hereby declare that this dissertation has not been submitted, in whole or in part, for any other degree, diploma or qualification at any other University. This dissertation is the result of my own work and includes nothing that is the outcome of work done in collaboration. I have attempted to reference appropriately any idea or finding that is not my own.

This dissertation does not exceed the limit of length specified by the Degree Committee for Biology, as stated in the Memorandum to Graduate Students.

Abstract

This thesis investigated the role played by regions of the prefrontal cortex and ventral striatum in the control of rats' behaviour by Pavlovian conditioned stimuli, and in their capacity to choose delayed reinforcement.

First, the function of the anterior cingulate cortex (ACC) in simple Pavlovian conditioning tasks was addressed. The ACC is a subdivision of prefrontal cortex that has previously been suggested to be critical for the formation of stimulus–reward associations. It was found that lesions of the ACC did not prevent rats from learning a simple conditioned approach response to a conditioned stimulus (CS) predictive of food reward, or from utilizing that CS as a conditioned reinforcer subsequently. Additionally, these subjects successfully acquired a conditioned freezing response to a CS predicting footshock. However, the same animals were impaired at the acquisition of autoshaped behaviour, an impairment that has been demonstrated previously. An autoshaping deficit was also observed when lesions were made following training. The phenomenon of Pavlovian–instrumental transfer was intact in these subjects. The hypothesis was developed that the ACC is not critical for the formation of stimulus–reward associations *per se*, but is critical when multiple stimuli must be discriminated on the basis of their differential association with reward. In support of this hypothesis, animals with lesions of the ACC were impaired on a version of the conditioned approach task in which a second, neutral stimulus, perceptually similar to the CS, was added; the lesioned subjects exhibited reduced discrimination.

Second, the role of the nucleus accumbens (Acb) in Pavlovian–instrumental transfer was investigated. The nucleus accumbens core, together with a larger amygdalar–striatal network of which it is a component, has previously been shown to be necessary for the expression of 'simple' Pavlovian–instrumental transfer. Rats with lesions of the nucleus accumbens core (AcbC) and shell (AcbSh) were tested on a 'response-specific' Pavlovian–instrumental transfer task, in which a Pavlovian CS selectively enhances instrumental responding for the outcome with which the CS was originally paired. AcbC lesions impaired the response specificity of this effect, while AcbSh lesions abolished Pavlovian–instrumental transfer entirely. These results are consistent with some — but not all — previous results in suggesting that the shell provides 'vigour' and the core provides 'direction' for the potentiation of behaviour by Pavlovian CSs.

Third, an attempt was made to train rats on a task for assessing preference for delayed reinforcement, using the 'adjusting-delay' paradigm. It was not immediately apparent that the rats reacted to the contingencies operative in this task, and mathematical analysis of their behaviour was conducted to establish whether their behaviour was sensitive to the delay, and what 'molar' features of performance on this task could be explained by delay-independent processes.

Fourth, a different delayed reinforcement choice task was developed, modifying a previously published task in which the subject is repeatedly offered a choice, in discrete trials, of a small reward delivered immediately, and a large reward delivered after a delay, with the delays systematically varied by the experimenter. Rats were trained on versions of this task in which the large, delayed reinforcer was or was not explicitly signalled by a cue present during the delay. The behavioural basis of performance on this task was examined, and *d*-amphetamine, chlordiazepoxide, and α -flupenthixol were administered systemically. It was found that the effects of *d*-amphetamine depended on whether the delayed reinforcer was signalled or unsignalled, increasing preference for signalled delayed reinforcement at some doses, but decreasing preference for unsignalled delayed reinforcement. These results may resolve contradictions in the literature, and are suggested to reflect the known effect of amphetamine to potentiate responding for conditioned reinforcers.

Fifth, rats that had been trained on this task (with no explicit signals present during the delay) were given lesions of the ACC, AcbC, or medial prefrontal cortex (mPFC). ACC-lesioned rats were no different from sham-operated controls in their ability to choose a large, delayed reinforcer. Lesions of mPFC reduced the tendency of subjects to shift from one lever to the other during the course of a session, but mPFC-lesioned subjects responded normally to removal of the delays, suggesting a loss of stimulus control. However, rats with lesions of the AcbC were severely impaired on this task, preferring the small, immediate reward, even though they discriminated the reinforcers. Additionally, the effects of intra-Acb amphetamine were assessed using a different version of the delayed reinforcement choice task, and found to have slight but inconsistent effects to reduce preference for the delayed reinforcer, though this effect did not depend on whether the delayed reward was signalled or unsignalled. These results suggest that the AcbC contributes significantly to the rat's ability to choose a delayed reward, a finding that has important implications for the understanding of Acb function. It is suggested that dysfunction of the AcbC may be a key element in the pathology of impulsivity.

Acknowledgements

First, I would like to thank Barry Everitt, who has been the best supervisor I could have wished for. His support, guidance, and encouragement made it all possible. I would also like to thank Trevor Robbins for his enthusiastic and insightful advice and support.

Thanks to John Parkinson for showing me the true path, for alcohol, for the coffee mornings of science. Thanks to Jeremy Hall for all those discussions, for showing me the many uses of electricity, for dinners and Withnail. Thanks to Tony Dickinson for explaining things patiently, and to Nick Mackintosh, not only for making available the facilities of his department, but for introducing me to the world of animal learning theory several years ago.

Thanks to Caroline Morrison for being a true friend and for knowing All about histology. Thanks to Anastasia Christakou for all the dinner excursions. Thanks to Cella Olmstead for getting me started and to Tim Bussey and Lisa Saksida for advice. Thanks to Nathaniel Daw for intellectual input, for running rats, and for beers underground. Thanks to all those others who helped run rats and were fun to work with: Guillaume Lachenal, Hosnieh Djafari Marbini, Andy Toner, David Pennicott, and Lakmali Sugathapala.

Thanks to Helen Sweet-Gossage for help with histology and instruction in the ways of the fen, to Rachel Underwood for histology, and to Simon Howes for teaching me stereotaxic surgery and many facts about cars.

Thanks to those others in the lab who made it such a fun place to be and enlivened trips from it, including Mercedes Arroyo, Maria Pilla, Pat di Ciano, Dan Hutcheson, and Kerrie Thomas.

Thanks to all those who looked after the animals so well, especially Mark Pauley and Julie Gautrey. Thanks to John Bashford, Hubert Jackson, Ian Cannell, John Picton, and John Baldwin, for help with photomicrography, matters electrical, slide presentation, metalwork, and computer networking respectively. I would also like to thank those researchers who took time to answer a stranger's questions so thoroughly, in particular Chris Bradshaw, James Mazur, John Evenden, and David Howell.

I am eternally grateful to those who taught me years ago and set me on this path; in particular, to Ian Gordon, Clare Mackay, Jane Lloyd, Duncan Lord, David Mansfield (sadly deceased), Geoff Cooke, Brian Dallaway, Bill Wright, and Tim Gill.

Thanks to Mark Rawlings, Carsten Zatschler, Ellen Rawlinson, and Rowan Cruft for more than I could ever repay. Thanks to Catrin Thomas for being there and putting up with so much. Thanks to Rob Milner, Paul Milligan, and Charles Offer for being fun to live with, to Sonia Zakrzewski for drinks and geckos, and to everyone else in and around Cambridge who partied. Thanks to the hundreds of composers, musicians, authors, and film directors who entertained me over the last three years. Finally, thanks to Ann and John Cardinal for everything.

Publications

To date, the following publications have arisen in whole or in part from this thesis:

Abstracts

1. **Cardinal RN, Everitt BJ, Robbins TW** (1999). Amphetamine interacts with cue stimuli to affect preference for delayed reinforcement. *Behavioural Pharmacology* **10** (supplement 1): S15–S16. (First Congress of the Behavioral Pharmacology Society and European Behavioural Pharmacology Society, 1–4 September 1999, Boston, Massachusetts, USA.)
2. **Cardinal RN, Lachenal G, Parkinson JA, Robbins TW, Everitt BJ** (2000). Effects of anterior cingulate cortex lesions on responding for conditioned reinforcement, discrete fear conditioning, autoshaping performance and Pavlovian–instrumental transfer. *European Journal of Neuroscience* **12** (supplement 11): 88 (abstract 44.8). (Federation of European Neuroscience Societies Second Forum Meeting, 24–28 June 2000, Brighton, UK.)
3. **Cardinal RN, Parkinson JA, Robbins TW, Dickinson A, Everitt BJ** (2000). Effects of lesions of the nucleus accumbens core and shell on response-specific Pavlovian–instrumental transfer. *Journal of Psychopharmacology* **14**(3) (supplement): A68 (abstract PH20). (British Association for Psychopharmacology Summer Meeting, 16–19 July 2000, Cambridge, UK.)
4. **Cardinal RN, Parkinson JA, Djafari Marbini H, Toner AJ, Robbins TW, Everitt BJ** (2000). Role of the anterior cingulate cortex in the control over behaviour by Pavlovian conditioned stimuli in rats. *Society for Neuroscience Abstracts* **26**: 980. (Society for Neuroscience 30th Annual Meeting, 4–9 November 2000, New Orleans, Louisiana, USA; abstract #366.13.)
5. **Everitt BJ, Parkinson JA, Lachenal G, Halkerston KM, Rudarakanchana N, Cardinal RN, Hall J, Morrison CH, Dalley JW, Howes SR, Robbins TW** (2000). Effects of limbic corticostriatal lesions on autoshaping performance in rats. *Society for Neuroscience Abstracts* **26**: 979. (SFN New Orleans; abstract #366.12.)

Papers and book chapters

1. **Everitt BJ, Cardinal RN, Hall J, Parkinson JA, Robbins TW** (2000). Differential involvement of amygdala subsystems in appetitive conditioning and drug addiction. In Aggleton JP (ed.), *The amygdala: a functional analysis* (second edition). Oxford University Press, New York (ISBN 0198505019), pp. 353–390.
2. **Parkinson JA, Cardinal RN, Everitt BJ** (2000). Limbic cortical-ventral striatal systems underlying appetitive conditioning. *Progress in Brain Research* **126**: 263–285. Chapter 17 of Uylings HBM, van Eden CG, de Bruin JPC, Feenstra MGP, Pennartz CMA (eds), *Cognition, emotion and autonomic responses: The integrative role of prefrontal cortex and limbic structures*. Elsevier, Amsterdam (ISBN 0444503323).
3. **Cardinal RN, Robbins TW, Everitt BJ** (2000). The effects of *d*-amphetamine, chlordiazepoxide, alpha-flupenthixol and behavioural manipulations on choice of signalled and unsignalled delayed reinforcement in rats. *Psychopharmacology* **152**: 362–375. (Available online via digital object identifier at <<http://dx.doi.org/10.1007/s002130000536>> or at <<http://link.springer.de/>>.)

Abbreviations used in this thesis

1. Pharmacology and chemistry

5-HIAA	5-hydroxyindoleacetic acid
5-HT	5-hydroxytryptamine (serotonin)
6-OHDA	6-hydroxydopamine
CCK	cholecystokinin
DA	dopamine
GABA	γ -aminobutyric acid
LiCl	lithium chloride
NMDA	<i>N</i> -methyl-D-aspartate
PB	phosphate buffer
PBS	phosphate-buffered saline

2. Psychology

CPP	conditioned place preference
CR	conditioned response
CRf	conditioned reinforcer
CS	conditioned stimulus
CS-	negative conditioned stimulus (i.e. predicts absence of stimulus)
CS+	positive, or predictive conditioned stimulus (be it appetitive or aversive)
CS ₀	neutral conditioned stimulus (strictly uncorrelated with stimulus)
CVD	conditional visual discrimination
ext	extinction
FI	fixed interval
FR	fixed ratio
IRI	interreinforcement interval
IRT	interresponse time
ISI	interstimulus interval
ITI	intertrial interval
NCRf	not-conditioned-reinforcer (e.g. a lever that produces no response; a control)
NS	not significant
OR	orienting response
PIT	Pavlovian to instrumental transfer
RI	random interval
RT	random time
S	subject (in analysis of variance notation); stimulus
UR	unconditioned response
US	unconditioned stimulus
VI	variable interval
VR	variable ratio

3. *Neurobiology and medicine*

ADHD	attention deficit/hyperactivity disorder
APCR	amphetamine potentiation of conditioned reinforcement
fMRI	functional magnetic resonance imaging
i.c.v.	intracerebroventricular
i.p.	intraperitoneal
i.v.	intravenous
IVSA	intravenous self-administration
MRI	magnetic resonance imaging (= nuclear magnetic resonance, NMR)
NMR	nictitating membrane reflex (eye-blink), typically in the rabbit
OCD	obsessive–compulsive disorder
PET	positron emission tomography
EEG	electroencephalogram
ERN	error-related negativity

4. *Neuroanatomy*

Abbreviations are those used by Paxinos & Watson (1996), except for those marked (*).

Acb	nucleus accumbens
AcbC	nucleus accumbens, core
AcbSh	nucleus accumbens, shell
ACC	anterior cingulate cortex (*)
AV	anteroventral nucleus of the thalamus
BLA	basolateral amygdala (BL) (*)
CeA	central amygdaloid nucleus (Ce) (*)
Cg1	cingulate cortex, area 1
Cg2	cingulate cortex, area 2
CPu	caudate putamen (striatum)
DLPFC	dorsolateral prefrontal cortex (*)
DS	dorsal striatum (*)
LC	locus coeruleus
LH	lateral hypothalamic area
MD	mediodorsal nucleus of the thalamus
NBM	nucleus basalis magnocellularis (nucleus of Meynert) (*)
OFC	orbitofrontal cortex (*)
OMPFC	orbitomedial prefrontal cortex (*)
PAG	periaqueductal grey
PCC	posterior cingulate cortex (*)
PFC	prefrontal cortex (*)
PRh	perirhinal cortex
SN	substantia nigra
SNc	substantia nigra pars compacta
SNr	substantia nigra pars reticulata
Sub	subiculum (S) (*)
VP	ventral pallidum
VS	ventral striatum (comprising Acb, ventromedial portions of caudate/putamen, olfactory tubercle) (*)

VSub	ventral subiculum (*)
VTA	ventral tegmental area

5. Physical units

These may be prefixed with k (kilo-, 10^3); c (centi-, 10^{-2}); m (milli-, 10^{-3}); μ (micro-, 10^{-6}).

A	amp ($C s^{-1}$)
C	coulomb
cd	candela
g	gram
h	hour (3600 s)
Hz	Hertz (s^{-1})
l	litre ($10^{-3} m^3$)
m	metre
M	molar ($mol l^{-1}$)
min	minute (60 s)
mol	mole ($\approx 6.022 \times 10^{23}$)
N	Newton ($kg m s^{-2}$)
s	second
W	watt ($kg m^2 s^{-3}$)

6. Statistics and probability

*	significant at $\alpha = 0.05$
**	significant at $\alpha = 0.01$
***	significant at $\alpha = 0.001$
ANOVA	analysis of variance
F	F statistic: the ratio of $MS_{\text{treatment}}$ to MS_{error}
MS	mean square
$P(X / Y)$	the probability of X, given that Y is true
P, p	probability
SD	standard deviation
SE	standard error
SED	standard error of the difference (between means)
SEM	standard error of the mean
α	threshold for determining statistical significance
$\tilde{\epsilon}$	Huynh–Feldt epsilon
μ	mean

7. General

\varnothing	diameter
$\sqrt{\quad}$	square root
W, D, H	(when referring to dimensions) width, depth, height
p. / pp.	page / pages

Chapter 1.

Introduction

OVERVIEW

This thesis investigates the role played by regions of the prefrontal cortex and ventral striatum in the control of rats' behaviour by Pavlovian conditioned stimuli, and in their capacity to choose delayed reinforcement.

In this introduction, I will first consider the ways in which animals, particularly rats, represent in their brains relationships or associations between environmental stimuli and motor acts. I will briefly review the evidence that Pavlovian conditioning is a distinct form of learning, and consider the types of associative representations that might be formed during Pavlovian conditioning. Similarly, I will attempt to provide a brief overview of instrumental conditioning, reviewing what is known about the associative representations that influence behavioural responses in the rat, the contribution of Pavlovian learning to instrumental performance, and the complex phenomenon of conditioned reinforcement.

Nearly a century's experiments on normal animals have given us insight into the kinds of associations that form in the *minds* of rats, and it is to be hoped that an accurate knowledge of these associations will assist in trying to understand the way in which rats' *brains* represent and control their world. Although some processes known to psychology may arise from the concerted action of many parts of the brain (e.g. consciousness; Baars, 1988), and some may occur ubiquitously across the nervous system (e.g. memory; Fuster, 1995), neuroscience has achieved considerable successes in mapping some psychological functions (e.g. specific sensory perception, initiation of movement) to distinct anatomical regions or chemical systems in the brain, sometimes fractionating psychological concepts in the process. With this in mind, I will outline the progress that has been made in mapping some of the neural systems responsible for the representations formed during Pavlovian and instrumental conditioning, focusing on the anterior cingulate cortex and the limbic corticostriatal circuitry of which it is a part, and using the perspective of the learning theories described.

Finally, I will consider another complex behavioural phenomenon: delayed reinforcement. Despite its functional importance, the mechanism by which delayed reinforcers control behaviour is not well understood psychologically or neurally, though tantalizing clues have been discovered. I will discuss the relationship between conditioned and delayed reinforcement, and the progress that has been made towards understanding their neural basis.

PSYCHOLOGICAL MECHANISMS FOR ACTION

Unlearned behaviour

Most basic among the mechanisms by which vertebrates influence the world, simple spinal and brainstem reflexes are critical for survival. Such reflexes influence skeletal musculature (respiratory movements, postural reflexes, pain-withdrawal reflexes, and the like) and autonomic function (such as the regulation of heart rate and arteriolar smooth muscle tone to maintain arterial blood pressure). Swallowing is a more complicated example of unlearned behaviour, involving the activation of at least ten different muscles in

a precisely-defined temporal order (Doty & Bosma, 1956). Indeed, innate behavioural patterns can be very complex: an oft-cited example is that of the female greylag goose, which exhibits an innate, species-specific and highly stereotyped behaviour (a ‘fixed action pattern’) in which it rolls eggs (or any vaguely similar object) into its nest; it will continue the movement even if the egg is lost or removed by an experimenter (Lorenz, 1939; Tinbergen, 1948).

Pavlovian conditioning

Pavlovian conditioning, or classical conditioning, is a term that refers to a set of experimental procedures, in which an experimenter arranges a contingency between stimuli in the world by presenting those stimuli independent of an animal’s behaviour. The term makes no assumptions about what is learned.¹ When an initially neutral stimulus (such as a bell) is paired with a unconditioned stimulus (US) (such as food) that elicits a reflexive or unconditioned response (UR), in this case salivation, then the bell becomes a conditioned stimulus (CS) that is now capable of evoking salivation as a conditioned response (CR).

Pavlov, the discoverer of this phenomenon (Pavlov, 1927), argued that a conditioned reflex developed because an association had formed between a representation of the CS and one of the US — stimulus substitution theory (Pavlov, 1927; Tolman, 1934). This would allow novel stimuli, through associative pairing, to control relevant species-specific response mechanisms, extending the usefulness of these responses. However, subsequent behavioural evidence has required the development of theories of conditioning that assume several associative representations are formed during the conditioning process. This evidence is summarized next.

Representations formed during Pavlovian conditioning

Associations are generally believed to be represented in the brain by altering the ‘weights’ of unidirectional synapses. As synaptic weights can only change on the basis of information available to the neurons involved (the ‘locality constraint’), the association of representations A and B in Figure 1 can only occur at points where information about these two representations converge, no matter what mechanisms exist to supervise and use the association. Such associations may be used for different purposes — for example, as a representation of a higher-order property of stimuli (a ‘feature detector’), or for commanding behavioural responses directly.




Figure 1. Simple associative representations. Plastic synapses change their weight on the basis of locally available information, requiring convergence of information about the representations to be associated.

In Pavlovian conditioning, there is the potential for several associations to form, as illustrated in Figure 2. In experiments in which the brain is damaged, lesions that removed a representation of either the CS, the US or the response would have obvious consequences not only for conditioned, but also unconditioned, responding (sites A, B, D in the figure). Lesions of site C, representing a central motivational state (such as fear) might not impair primitive unconditioned responses, yet could impair conditioned responses that

¹ Strictly, ‘classical conditioning’ is the generic term and refers to the experimental procedure described, while ‘Pavlovian conditioning’ implies Pavlov’s interpretation of the process (Mackintosh, 1974, pp. 96/98). This distinction will not be followed in the present discussion, and the term ‘Pavlovian conditioning’ will be used without implying an underlying mechanism.

were based on the elicitation of fear. Again, however, any properties of the unconditioned response to the US that depended on this hypothesized ‘fear’ state would be lost.

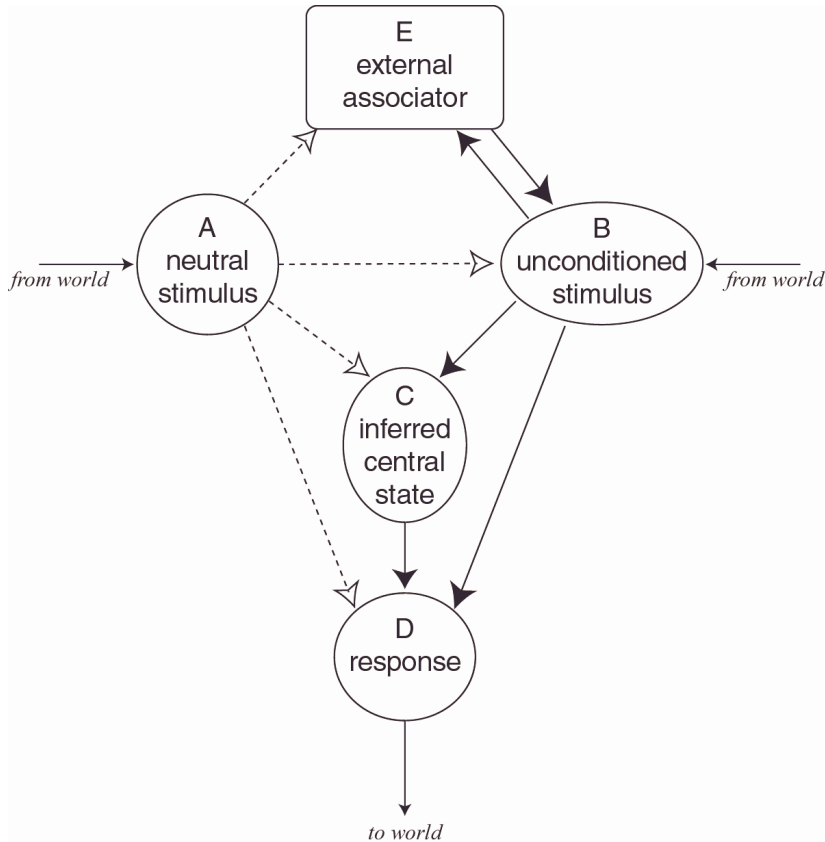


Figure 2. From a theoretical perspective, Pavlovian conditioning has the potential to create associations between a conditioned stimulus (CS) and representations of the unconditioned stimulus (US), central states such as fear, and unconditioned responses. Only a single response is shown; distinctions between different kinds of response are discussed in the text. Dotted lines represent associative links. Bidirectional communication also allows representations to be associated in ‘third-party’ sites (E) (for similar ideas, see Damasio & Damasio, 1993; Fuster, 1995, p. 88). Note that lesions of such a site might prevent conditioning without impairing any form of unconditioned response, as would selectively disconnecting the CS from a representation involved in responding.

Experimental analysis of Pavlovian conditioning has shown that CS–US pairings may cause the CS to elicit at least three of these representations in the brain (Dickinson, 1980; Mackintosh, 1983; Gewirtz & Davis, 1998). The first and simplest of these is that the CS may become directly associated with the *unconditioned response* (UR), a simple stimulus–response association.

The second is a representation of *affect* — such as fear or the expectation of reward — as demonstrated by the phenomenon of transreinforcer blocking, in which the presence of a CS previously paired with shock can block or prevent conditioning to a CS paired with the absence of otherwise expected food reward (Dickinson & Dearing, 1979). These two reinforcers share no common properties other than their aversiveness and therefore the blocking effect (see Kamin, 1968; 1969) must depend upon an association between the CS and affect. Affective states can therefore be independent of the specific reinforcer and response. This concept has been widely used in theories of learning (Konorski, 1948; Konorski, 1967; Dickinson & Dearing, 1979) and is illustrated in Figure 3. Associations between the stimulus and an affective state appear to be critical in second-order conditioning (in which $S_1 \rightarrow US$ pairings are followed by $S_2 \rightarrow S_1$ pairings); unlike a first-order conditioned response (CR), a second-order CR is relatively insensitive to post-training changes in the value of the US (implying that the second-order CR does not depend on $S_2 \rightarrow US$ associations) and the response to S_2 may differ from the response to S_1 or the US (implying that it does not depend on $S_2 \rightarrow R$ associations) (reviewed by Gewirtz & Davis, 1998).

The third form of representation is *specific to the US*. If a CS is paired with a desirable food and the food is subsequently devalued by pairing it with LiCl injection to induce nausea, so that the food becomes aversive and is rejected, the reaction to the first-order CS changes in normal animals (Mackintosh, 1983).

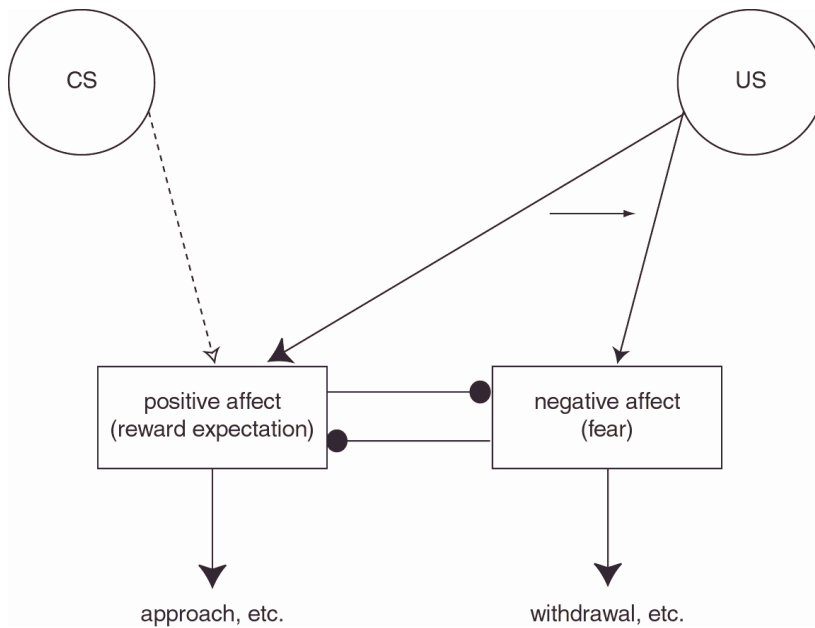


Figure 3. Conditioning to affective states leaves the response independent of the current value of the US. (The CS–affect association corresponds to link A→C in Figure 2.) The CS associates with the affective state elicited by the US during conditioning, but if the US subsequently alters its value (shown as a shift in the US–affect link), the conditioned response (CR) will not alter. The links between the affective states are inhibitory, reflecting the supposition that appetitive and aversive affective or motivational states are mutually antagonistic (Konorski, 1967; Dickinson & Dearing, 1979) (see Mackintosh, 1983, pp. 114–123), though this issue is not important for the present discussion.

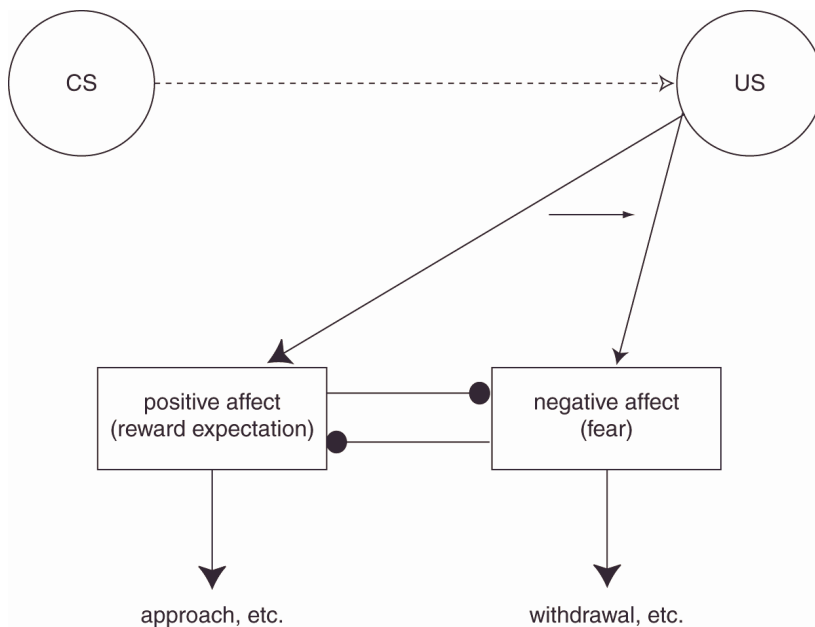


Figure 4. CS–US conditioning allows the conditioned response to alter and reflect changes in the unconditioned response induced by US devaluation. (The CS–US association corresponds to link A→B in Figure 2.) Compare Figure 3.

Therefore the CS cannot have been associated with an abstract ‘positive affect’ representation, but must have been associated with that particular reinforcer. The association must be specific to the US, because the reaction to a second CS that predicted a different food does not alter, and its connections with valence information must be modifiable and *downstream* of the CS–US association (Figure 4).

Further evidence that a CS becomes associated with a relatively specific representation of the properties of a reinforcer is provided by studies of cued instrumental discrimination (Trapold, 1970). Rats acquired an instrumental discrimination between two levers paired with different appetitive reinforcers more rapidly if the discriminative cues had been paired with the same reinforcers (same condition) in a previous Pavlovian stage than when the outcome was switched between stages (different condition). A rigorous demonstration of the formation of associations between the *specific sensory properties* of stimuli comes from sensory preconditioning (Brogden, 1939), the process by which neutral stimuli are paired in the form $S_2 \rightarrow S_1$, after which $S_1 \rightarrow US$ conditioning causes a CR to occur to S_2 . Thus, in the first stage, associations form between representations that have no motivational component. Taken together, these pro-

cedures demonstrate that animals are able to encode the relationship between a CS and specific sensory properties of a US and furthermore that they can relate this sensory representation to the affective valence of the reinforcer.

It is likely that further representations are formed during Pavlovian conditioning; for example, it has been argued that animals remember the precise intervals between CS and US presentation, and even that this process — rather than simple associative learning — determines conditioned responding (Gallistel, 1994), though these issues will not be discussed further.

Goal-directed behaviour

The term ‘instrumental conditioning’ refers to an experimental procedure in which the experimenter arranges a contingency between an animal’s behaviour and a reinforcing outcome (Thorndike, 1911). No assumptions are made about the nature of learning; what an animal does in fact learn has been a matter of debate for decades. Although arguments were once made that instrumental conditioning was explicable in terms of Pavlovian conditioning, and vice versa, the two have been doubly dissociated. Classical conditioning is not explicable purely in terms of instrumental response–reinforcer contingencies (for example, Sheffield, 1965, demonstrated conditioned responding even though responding led to the omission of reinforcement). Similarly, instrumental conditioning cannot be explained entirely in terms of classical conditioning; one demonstration was that of Grindley (1932), who employed a form of *bidirectional control* assay. He trained guinea pigs to turn their heads to one side in order to earn a piece of carrot delivered immediately in front of them, and then altered the instrumental contingency, training the same subjects to turn the other way. As there is no differential Pavlovian contingency between the two responses, the asymmetry in responding must have been due to the instrumental contingency. What, though, is learned as a result of this contingency?

Early theorists took the position that the delivery of reward strengthened an associative connection between environmental stimuli and a particular response (Thorndike, 1911; Grindley, 1932; Guthrie, 1935; Hull, 1943). Such learning would represent *procedural* knowledge (Dickinson, 1980), as the structure of the representation directly reflects the use to which the knowledge will be put in controlling the animal’s behaviour, and would be inflexible, in that subsequent changes in the value of the reward would be unable to affect responding.

However, it has been shown that rats form more sophisticated and flexible representations in instrumental conditioning tasks. Behaviour may be said to be *goal-directed* if it depends on the twin representations of (1) the instrumental contingency between an action and a particular outcome, and (2) a representation of the outcome as a goal (Tolman, 1932; Dickinson & Balleine, 1994). Simply put, a goal-directed organism presses a lever for food because it knows that lever-pressing produces food *and* that it wants the food. As performance of such behaviour requires these two representations to interact, the knowledge upon which performance is based must be *declarative* — that is, the knowledge is to some degree independent of the use to which it is put. (The interaction is dynamic, and if two actions are to be compared simultaneously, a binding problem occurs — if action 1 is associated with outcome 1, which is associated with high incentive value, that high value must be ‘mapped back’ to action 1, not action 2 or 3, to encourage its performance. The problem of representing such symbolic processing in a connectionist network like the brain is discussed by Holyoak & Spellman, 1993; Shastri & Ajjanagadde, 1993; Sougné, 1998.)

Instrumental contingency

Free-operant lever-pressing in rats satisfies the instrumental criterion, as shown by a bidirectional control assay (Bolles *et al.*, 1980). Not all behaviour may be conditioned instrumentally, however; for example, scratching is difficult to condition in rodents (Shettleworth, 1975; Morgan & Nicholas, 1979). Similarly, the instrumental status of spatially directed locomotion is in doubt, as illustrated by Hershberger (1986) using a ‘looking-glass runway’ in which chicks were required to run away from food in order to obtain it; the chicks were unable to do so.

Incentive value

The status of an outcome as a goal may be determined by an outcome devaluation test, so long as the test controls for Pavlovian conditioning and is conducted in extinction, so there is no chance for subjects to alter their behaviour by learning a new action–outcome relationship (Dickinson & Balleine, 1994). Using such a procedure, Adams & Dickinson (1981) demonstrated the goal status of food in lever-pressing tasks using hungry rats. The rats were given access to a lever, and one type of food (termed the positive outcome) was delivered contingent upon lever-pressing. Another type of food (termed the negative outcome) was delivered if the rats refrained from responding for a short while. Adams & Dickinson then devalued the positive outcome for one group and the negative outcome for another by pairing it with LiCl injection, and tested the rats in extinction (so there was no opportunity to learn a new response–outcome relationship). It was found that the rats for which the positive outcome was devalued pressed the lever less than those rats for which the negative outcome was devalued. This result cannot be attributed to a general suppressive effect of the devaluation procedure on lever-pressing. As Pavlovian conditioning was also controlled for, the differential effect of devaluing the ‘positive’ and ‘negative’ outcomes must have been mediated by the instrumental contingency, a result confirmed by Colwill & Rescorla (1985) using a choice procedure.

However, under certain circumstances, the goal status of the food does not alter *immediately*. For example, if the food is devalued by isotonic LiCl injection following a meal, rats do not work less for the food until they have had the opportunity to *re-experience* the food by consuming it (Balleine & Dickinson, 1991). This implies that there are two representations of the food’s value. After the injection, something in the rat has learned from the conditioning experience, and will react ‘aversively’ to the food when it is next eaten, but the *instrumental incentive value* (the value governing goal-directed instrumental action) has not yet changed. Dickinson and colleagues refer to the process by which instrumental incentive value is updated as *incentive learning* (Dickinson & Balleine, 1994).

Hedonic assessment: an implied, distinct valuation process

The system that reacts immediately (but covertly) to food devaluation procedures, is independent of the instrumental incentive value, and comes into play upon direct experience of the food has been termed an affective or *hedonic* system (Garcia, 1989). To restate this hypothesis: the devaluation procedure modifies the neural system responsible for hedonic experience, so that it will react with disgust rather than pleasure when the devalued foodstuff is next experienced. In the meantime, the more ‘cognitive’ incentive value remains high, so the animal still works for the devalued food. The next time the food is consumed, direct experience of the food leads to the disgust reaction being evoked, which re-writes the neural representation of incentive value and leads the animal to work less for the food in the future.

Although hedonic reactions may be conditioned and assessed in humans by direct questioning (e.g. Baeyens *et al.*, 1990), it is not obvious that they can be assessed at all in other species. However, it has been suggested that taste reactivity patterns — the orofacial reactions of rodents to flavours introduced

into the mouth — are an index of hedonic experience in rats (Grill & Berridge, 1985), and indeed, they behave in a manner compatible with the role required by Dickinson and colleagues of their hedonic system, such as tracking motivational state directly (Berridge *et al.*, 1981; Berridge *et al.*, 1984; Berridge, 1991; Berridge & Robinson, 1998, p. 314).

Some other features of the incentive learning process are worth noting. Some treatments, such as hypertonic intraperitoneal LiCl injection (which causes somatic discomfort) change both ‘value’ representations at the same time (Balleine & Dickinson, 1992). Additionally, the incentive value of a food is initially independent of current motivational state (i.e. hunger, thirst); thus, when a hungry rat is trained to respond for food, and then sated before being tested in extinction, it will respond as much as another rat that is hungry on test, until it experiences directly the reduced value of the foodstuff when sated (Balleine, 1992). Following this re-experience, the incentive value will vary appropriately with the motivational state of the animal, so that the rat will subsequently work hard if it is hungry, but not if it is sated. This implies that both the ‘immediate-assessment’ (hedonic) system and the incentive value system have access to motivational state information.

Discriminative stimuli

A very brief review of the role of discriminative stimuli (S^D) is also in order. The basic procedure for establishing a stimulus as a positive S^D is to reinforce responding during the S^D , and to withhold reinforcement in its absence. Clearly, while such a stimulus does signal that the instrumental response–reinforcer contingency is in operation, and might ‘set the occasion’ for responding, it might also enter into direct stimulus–response associations, or act as a Pavlovian CS for the reinforcer (because the reinforcer is only delivered in the presence of the S^D). Indeed, it has been demonstrated that S^D s do become associated with their reinforcer (reviewed in Colwill & Rescorla, 1988). Nevertheless, it is likely that a more complex explanation is also required: S^D s have effects that cannot be explained in terms of Pavlovian associations (Holman & Mackintosh, 1981) and it has been shown there is a conditional relationship in which an S^D signals the operation of a particular response–reinforcer contingency (Colwill & Rescorla, 1990; Rescorla, 1990a; Rescorla, 1990c).

Stimulus-response habits

Although rats possess declarative knowledge of the consequences of their actions, this does not mean that they lack a procedural stimulus–response ‘habit’ system. There have been a number of demonstrations in which reinforcer devaluation failed to affect instrumental responding (reviewed by Adams, 1982). Investigating the reason for these findings, Adams (1982) established that overtraining is one critical determinant of whether an instrumental response becomes ‘autonomous’ and resistant to devaluation. Following limited experience of instrumental training (such as training in which 100 reinforcers were earned), rats’ actions remained under the control of the instrumental contingency, and were responsive to reinforcer devaluation. With extended experience of instrumental responding (such as training in which 500 reinforcers were earned), their actions became habitual and resistant to devaluation (see also Dickinson *et al.*, 1995). With ratio schedules, the number of reinforcers is more relevant than the number of times the action is practised (Adams, 1982).

The schedule of reinforcement is one other factor that has an important influence on habit development. Actions trained on interval schedules are more likely to become habitual than those trained on ratio schedules (Dickinson *et al.*, 1983), presumably because of the weaker response–reinforcer contingency that such a schedule involves (Dickinson, 1994). It has been argued that a low level of experience of this contingency is the central factor governing habit development (Dickinson, 1985).

Pavlovian to instrumental transfer

Pavlovian stimuli can modulate instrumental performance by at least two mechanisms (Dickinson, 1994; Dickinson & Balleine, 1994). For example, a stimulus that predicts the arrival of sucrose solution will enhance lever-pressing for sucrose; this is the basic phenomenon of Pavlovian-to-instrumental transfer (PIT) (Estes, 1948; Lovibond, 1983). These stimuli may have a *general* motivating effect, so that a CS for a sucrose solution will enhance lever-pressing for sucrose — but also for dry food pellets — when the animal is thirsty. In addition, they may act selectively to potentiate actions with which they share an outcome (in this example, potentiating lever-pressing for sucrose but not for food); this is a *response- or outcome-specific* form of PIT.

Response-specific potentiation

A prototypical demonstration of this effect was provided by Colwill & Rescorla (1988, Experiment 3, abbreviated). Animals were trained to associate a stimulus with either pellets or sucrose solution. They were then trained separately to perform two instrumental actions, one for pellets and one for sucrose. The design is shown in Table 1.

Table 1. Specific Pavlovian–instrumental transfer; design of Colwill & Rescorla (1988) (S, stimulus).

Group	Training	Test
Pellet	S → pellet Lever-press (Lp) → pellet Chain-pull (Cp) → sucrose	S: Lp > Cp
Sucrose	S → sucrose Lp → pellet Cp → sucrose	S: Lp < Cp

Animals are hungry throughout.

During an extinction test, the stimulus had a greater ability to potentiate the action with which it shared an outcome (see also Colwill & Motzkin, 1994). (Note that response specificity is implied whether the stimulus selectively elevated the action with which it had shared an outcome, or selectively depressed the action with which it had *not* shared an outcome, as Colwill & Rescorla actually observed.) It has been suggested that Pavlovian stimuli act by reinstating conditions that are more similar to those in which the instrumental action was trained (Trapold & Overmier, 1972), but this need not be the case. Colwill & Rescorla (1988) also compared the effects of a Pavlovian CS+ to those of a discriminative stimulus (S^D); though the transfer effect was qualitatively similar there were quantitative differences and they suggest differences between the roles of a CS and an S^D. The exact mechanism is therefore controversial (see Colwill & Motzkin, 1994), but the concept of response-specific PIT has been clearly demonstrated.

It should be noted that the potentiation is dependent upon the instrumental response and outcome, and necessarily dependent upon the Pavlovian US. The CS must call up a representation of the US sufficiently detailed to discriminate pellets from sucrose.

General potentiation of behaviour

Dickinson & Dawson (1987b, Experiment 2) showed that Pavlovian stimuli whose outcomes are relevant to the current motivational state may potentiate ongoing instrumental behaviour in a general fashion. Their design is illustrated in Table 2.

Table 2. General PIT; design of Dickinson & Dawson (1987b).

Group	Training	Test
Hungry	S1 → pellet	(split into two groups)
	S2 → 20% sucrose	Hungry: $Lp(S1) > Lp(S2)$
	Lp → pellet	Thirsty: $Lp(S1) < Lp(S2)$

Following training, all animals were tested in extinction. Half were tested hungry, and in this case presentation of the stimulus associated with pellets had the greatest effect to increase lever-pressing. However, half were tested thirsty; here, the liquid sucrose solution is more relevant to the current motivational state, and indeed the animals pressed more under the stimulus associated with sucrose — despite never having pressed a lever for sucrose. The potentiation is thus independent of the outcome of the instrumental action. It may be clearly differentiated from PIT based on the reinstatement of training conditions, for such an effect would predict $Lp(S1) > Lp(S2)$ in all cases.

It may be argued that the thirsty rats in this study are showing specific suppression by S1, rather than general potentiation by S2. Indeed, there is evidence that stimuli predictive of food suppress lever-pressing in thirsty rats. A study by Balleine (1994) demonstrated such an effect, which was asymmetrical across hunger and thirst and which may reflect the fact that dry foods aggravate water deprivation and are thus aversive to thirsty rats, while fluid consumption does not aggravate food deprivation. However, this same study also provided clear evidence of a general PIT effect. Rats that were trained to lever-press for water whilst thirsty, and were then shifted to a state of hunger, pressed more when a CS for food pellets was presented (Fig. 2 of Balleine, 1994, group PEL, tested hungry). A general suppressive effect was also seen when a CS for food was presented to thirsty rats previously trained to press for water.

This general potentiation has been interpreted as *conditioned motivation* (see Dickinson, 1994): motivation is conditioned to the Pavlovian CS during training. Note that the potentiation is independent of the instrumental outcome (a stimulus paired with sucrose solution can potentiate responding for pellets), and obviously independent of the instrumental response. However, it is dependent upon the Pavlovian US because the CS only potentiates behaviour when its US is currently relevant, and therefore the US must be distinguished from others that are not.

This process is responsible for the ‘irrelevant incentive’ effect, in which stimuli (including contextual stimuli) associated with motivationally relevant outcomes potentiate behaviour, even in the absence of prior experience of the outcome in that motivational state (Dickinson, 1986; Dickinson & Dawson, 1987a; 1987b). That is, the level of Pavlovian–instrumental transfer depends directly on the relevance of the outcome to the current motivational state; therefore, the neural system responsible for PIT must have access to motivational state information. As the potentiation is independent of instrumental outcome, by implication it cannot affect choice behaviour; this has been demonstrated for the irrelevant incentive effect (Dickinson, 1986).

The implication is that information about the identity of the US must be available on test. In Dickinson & Dawson’s (1987b) study, it would not be sufficient to have learned affective values during training (S1 → nice, S2 → nice), even if they had slight differences (S1 → superb, S2 → OK) because this could not explain the bidirectional nature of the stimulus control on test. Furthermore, there is no opportunity to learn a conditional affective value (S1 + hunger → nice, S1 + thirst → nasty, etc.) because the reinforcer has never been experienced in the motivational state present on test. Therefore, a more detailed representation of the US must be available (such as its liquidity), which can then be assessed for motivational relevance at the moment of test. The degree of specificity can be gauged by comparing similar rein-

forcers; for example, this procedure requires the discrimination of sucrose solution and pellets during the test using only conditioned cues, and rats are clearly capable of this.

As specific and general PIT have never been doubly dissociated, it remains possible that they share a common mechanism — for example, that one process provides ‘vigour’, which is then ‘directed’ to appropriate responses if they are available. However, it seems likely that PIT is separable from the ‘hedonic’ system discussed above: for example, PIT (in the form of the irrelevant incentive effect) can occur without updating instrumental incentive values (Dickinson, 1986).

Simple transfer tasks

Pavlovian–instrumental transfer has more often been assessed in a simpler task. The general design is illustrated in Table 3.

Table 3. Simple designs for demonstrating PIT. (ISI, interstimulus interval.)

Study	Training	Test
Lovibond (1983)	S+ → pellet	Lp(S+) > Lp(ISI)
Estes (1948)	Lp → pellet	
Dickinson <i>et al.</i> (2000)	S+ → pellet	Lp(S+) > Lp(ISI)
Hall <i>et al.</i> (1999)	S- → Lp → pellet	Lp(S+) > Lp(S-)

Animals are hungry throughout.

This type of task does not differentiate between the two transfer mechanisms, as the stimulus shares a motivationally relevant outcome with the only action being tested.

Simple PIT was further investigated by Lovibond (1983), who found that Pavlovian CSs had different effects on responding under ratio and interval schedules. Under variable ratio (VR) schedules, PIT was only observed when the pre-CS response rates were very low, and the potentiation of responding was long-lasting. Essentially, the CS *restarted* subjects that had ceased responding. In contrast, when variable interval (VI) schedules were used, presentation of the CS resulted in an elevation of responding that lasted only for the duration of the CS, and the CS rate of responding was approximately proportional to the pre-CS rate of responding. This result may be interpreted in several ways. Lovibond, for example, argued that the CS was effective when there was least ‘stimulus support for responding’ (i.e. when temporal cues from the schedule or the animal’s own responding predicted food least; Lovibond, 1983). As interval schedules engender habitual responding more readily than ratio schedules (Dickinson, 1985; 1994), PIT may reflect a process that enhances the habitual component of instrumental responding. Alternatively, PIT may boost goal-directed (non-habitual) responding by acting as a form of reminder, thus restarting responding under ratio schedules at moments when the contribution of goal-directed responding was low, and consistently enhancing responding under interval schedules, when the contribution of goal-directed responding is always low (A. Dickinson, personal communication, 1999). It is also possible that general and response-specific PIT have different effects in this regard; all these questions remain unanswered. Regardless, interval schedules appear to provide the best conditions under which to observe a stable and reliable PIT effect.

Summary

At least six processes are implicated in instrumental performance in rats. This picture, undoubtedly complex, is summarized in Figure 5. Although some interactions between Pavlovian and instrumental behaviour have already been discussed, there remain many uncertainties, such as whether central ‘affect’ states

(inferred from Pavlovian studies) are the same as, or interact with, one of the ‘value’ systems inferred to control instrumental responding.

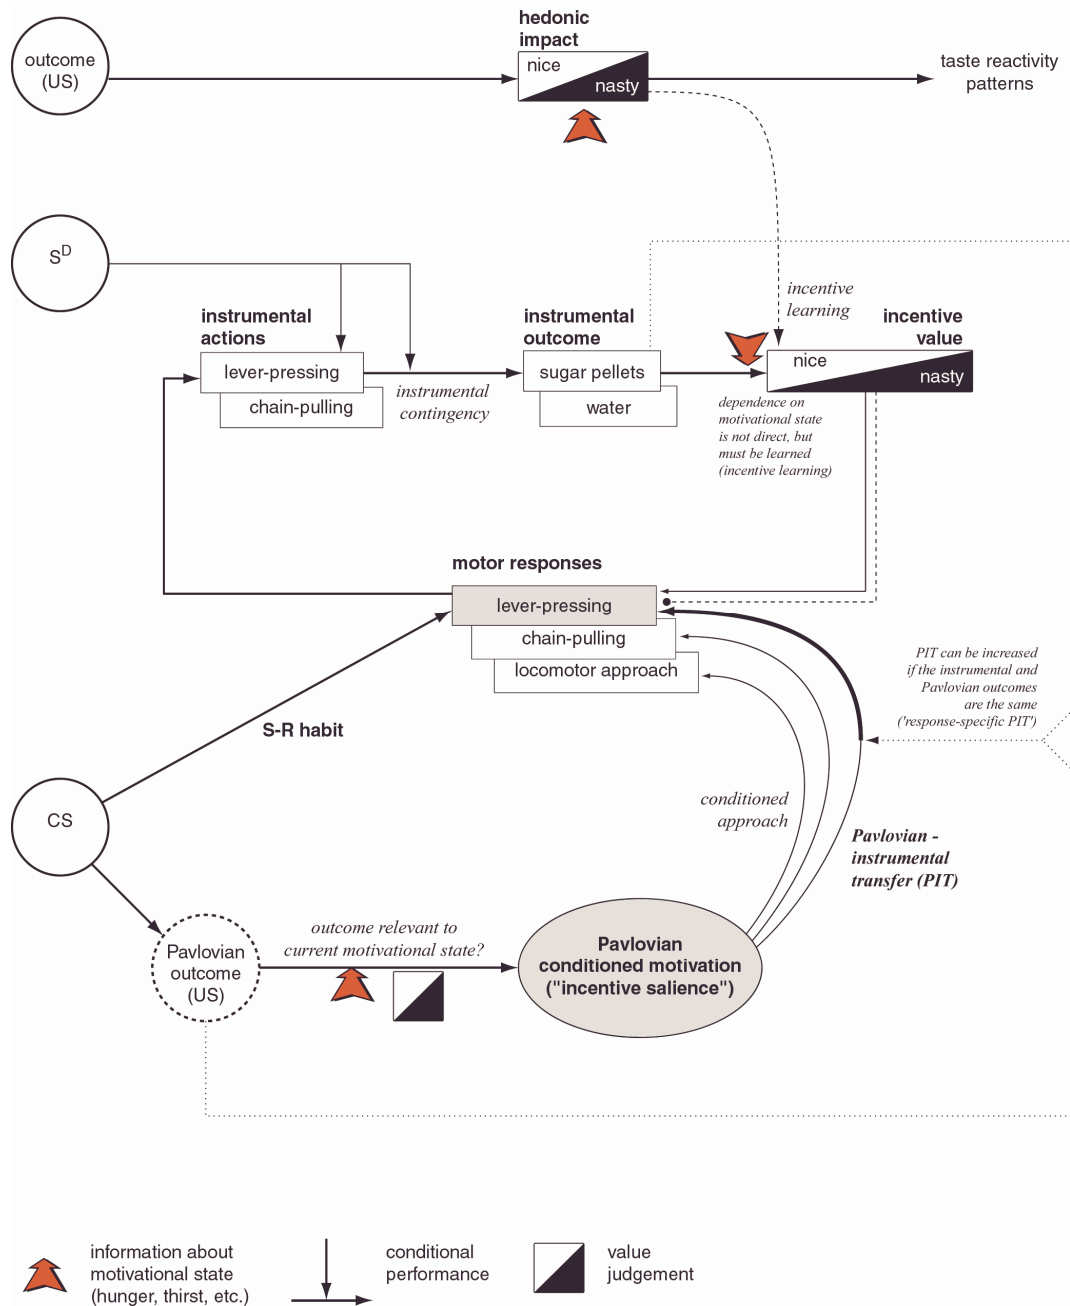


Figure 5. Some processes that contribute to instrumental behaviour in rats. An action such as lever-pressing is capable of being detected and represented in a system that can encode the contingency between this action and outcomes. When this representation is combined with a favourable representation of the instrumental incentive value of the outcome, lever-pressing is promoted. The instrumental contingencies currently in force can be signalled by instrumental discriminative stimuli (S^D). The value governing goal-directed responding is learned through direct experience of the outcome in particular motivational states; it can therefore be distinguished from a ‘hedonic’, or immediate-assessment value system (see text). A separate contribution to response output comes from direct stimulus–response associations (S–R habits), which can be formed through repeated training. In addition to these processes, Pavlovian conditioned stimuli (CSs) that signal a motivationally relevant outcome can enhance responding, both by providing a ‘motivational boost’ and by potentiating responses that share an outcome with the Pavlovian CS.

CONDITIONED REINFORCEMENT

When a CS is paired with reward, it may gain some of the properties of that reward and be capable of reinforcing behaviour itself; it is then termed a conditioned reinforcer (see Mackintosh, 1974, pp. 233–259). Stimuli are established as conditioned reinforcers by Pavlovian association, and affect instrumental performance, although it is not absolutely clear how conditioned reinforcement relates to the Pavlovian and instrumental processes described so far.

Validity of conditioned reinforcement as a concept

The concept that stimuli may themselves become reinforcing has not always been accepted (recently reviewed by Williams, 1994a). Though recent demonstrations have proven that conditioned reinforcement is a genuine phenomenon, it is sometimes difficult to interpret the role of putative conditioned reinforcers in behavioural experiments. There are two common reasons for this: one relates to the techniques used to *measure* conditioned reinforcement; and the other to the methods used to *establish* stimuli as conditioned reinforcers, as some of these methods may endow the stimuli with other functions.

It is frequently suggested that when a response-contingent stimulus predicts reward, the stimulus becomes established as a conditioned reinforcer (CRf). For example, rats can be trained under second-order schedules of reinforcement; in a typical schedule, denoted FR10(FR5:S), every fifth response produces a stimulus and every tenth stimulus is accompanied by reinforcement. In this situation, omission of the stimulus impairs performance (e.g. Arroyo *et al.*, 1998). Though the stimulus is certainly paired with reinforcement, it is clear that it could play a discriminative role in this task (signalling that responding is likely to lead to reinforcement); depression of responding when the stimulus is omitted would not be direct evidence that the stimulus has conditioned reinforcing properties.

However, conditioned reinforcers have effects beyond those of an S^D . It has been shown that performance under second-order schedules depends to some extent on the association of the stimulus with primary reinforcement (reviewed by Mackintosh, 1974, p. 241). More directly, Williams & Dunn (1991) demonstrated using a choice schedule that pigeons preferred a key associated with a CRf despite this leading to more unreinforced trials. In this study, preference tracked the frequency of extinction trials in which the conditioned reinforcer was presented, unconfounded by differences in primary reinforcement or by changes in the value of the CRf itself.

Further alternative explanations have been offered for the effects of response-contingent stimuli predictive of reward; for example, it has been suggested that they exert their effects through ‘marking’ the fact that a response has been made, or ‘bridging’ delays. Although these effects undoubtedly occur, it has been shown that conditioned reinforcers have *reinforcing* properties over and above their other functions (Williams, 1991).

When conditioned reinforcement must be demonstrated directly, the cleanest technique uses a different approach, namely the ‘acquisition of a new response’ procedure (Mackintosh, 1974, pp. 235–237). In an initial phase, a CS is paired with reward; in a second phase, the CS is provided contingently upon a response that is new to the subject. If changes in general activity are controlled for by providing two such new responses (i.e. a choice procedure) and the subjects work for the CS more than they work for an unpaired stimulus, it can be concluded that the stimulus is functioning as a conditioned reinforcer.

Relationship of conditioned reinforcement to other processes controlling instrumental performance

Although it seems reasonably clear that conditioned reinforcers acquire some of the properties of the primary reinforcer, rather than simply by providing information about its availability (Mackintosh, 1974, pp. 250–259), it is not immediately clear how to relate conditioned reinforcement (CRf) to the ‘value systems’ governing instrumental behaviour described earlier. It is relevant to consider the relationship, though, because there is considerable evidence regarding the neural basis of responding for conditioned reinforcement, which will allow comparison with the neural basis of other factors governing instrumental performance. Several possibilities exist; it should be stated at the outset that there is relatively little experimental evidence concerning any of them.

Reinforcement of S–R habits by the CRf. One simple possibility is that the CRf reinforces behaviour in the way that Thorndike (1911) originally envisaged reinforcement in his Law of Effect — by ‘stamping in’ stimulus–response associations. Were this to be the case, the resulting behaviour should be resistant to devaluation of the CRf (assessable, perhaps, by Pavlovian CRs evoked by the CRf). Although there have been no direct studies of this question, extinction of a CRf rapidly affects preference even when the contingency between responding and primary reinforcement is held constant (Dunn *et al.*, 1987), suggesting that responding for conditioned reinforcement is not always habitual. Whether it can *become* habitual may be a very difficult question to answer: establishment of habits with primary reinforcement requires considerable training (Adams, 1982), and if conditioned reinforcement is used with an ‘acquisition of new response’ procedure, the CRf–US relationship will extinguish as the response is being acquired.

General PIT. We have already seen that noncontingent CSs cannot affect *choice* behaviour via the general PIT mechanism, and a mechanism that does not affect choice could not be accepted as conditioned reinforcement. This would seem to rule out general PIT as an explanation for CRf. However, when the CS is *contingent* upon one response but not another, a general potentiation of behaviour could affect choice. Pressing a lever that produces a CRf lever effectively gains access to a stimulus that ‘boosts’ behaviour (albeit unselectively), while pressing a control (NCRf) lever does not. If the animal’s current behaviour is boosted, the CRf would tend to boost the behaviour that caused its presentation; this mechanism could favour CRf over NCRf responding. Of course, PIT could not affect responding before the first presentation of the CRf, so seems unlikely to account for results such as those of Williams & Dunn (1991), using concurrent-chain schedules (in which responding is measured for a period before the animal earns a CRf), but it is possible that PIT contributes to responding under free-operant schedules in which CRf presentation can occur during periods of responding and the CRf is presented for a comparatively long time.

Response-specific PIT. As the response-specific PIT effect increases responding for actions that share an outcome with the Pavlovian CS, it seems unlikely to affect responding for conditioned reinforcement, because the outcome of the response (which is the CRf) is not the same as the outcome of the CRf (which is the primary reinforcer). Potentiation could only occur through this mechanism if the brain could apply a less strict definition of outcome, thus: ‘lever-pressing shares an outcome with the CS because lever-pressing produces the CS and therefore (by a process of inference or association) the US’. This is akin to second-order conditioning, however, so the mechanism should not be ruled out.

Instrumental incentive value of the CRf. The final possibility is perhaps the most obvious — that the CRf becomes an instrumental goal, possessing incentive value itself, which it acquires through an affective or hedonic process. Direct demonstration of this would involve training animals to respond for a conditioned reinforcer, devaluing the CRf, and assessing responding in extinction (with appropriate controls;

cf. Adams & Dickinson, 1981). Nevertheless, even in the absence of such a demonstration, it seems very likely from the results of Williams and colleagues (Dunn *et al.*, 1987; Williams & Dunn, 1991) that conditioned reinforcers do acquire incentive value; these authors have demonstrated effects of conditioned reinforcers on choice using tasks to which PIT is unlikely to contribute and in which responding is probably not habitual.

In summary, though the evidence is far from conclusive, it seems most likely that conditioned reinforcers control behaviour by acquiring instrumental incentive value, though they might also act to potentiate responding further via PIT.

NEURAL DISSOCIATIONS WITHIN ASSOCIATIVE LEARNING

Overview of the limbic corticostriatal circuits considered in this thesis

While investigating hypothalamic function, Hetherington & Ramsay (1939) and Anand & Brobeck (1951) found that electrolytic lesions of the lateral hypothalamus (LH) appear to leave animals demotivated — with impairments in unlearned behaviour (subjects were aphagic and adipsic, with reduced sexual, exploratory, and maternal behaviour) and in learned behaviour (impaired instrumental responding). However, such lesions also disrupt the medial forebrain bundle, a fibre tract that passes through the lateral hypothalamus and contains the dopaminergic projection from midbrain dopamine (DA) neurons (the substantia nigra pars compacta, SNc, and the ventral tegmental area, VTA) to the forebrain. As lesions of this projection using the dopamine-depleting toxin 6-hydroxydopamine (6-OHDA) produced a similar pattern of behavioural impairment (Marshall & Teitelbaum, 1977), as did dopamine-depleting lesions of one its targets, the striatum (Stricker & Zigmond, 1976), attention was focused on the behavioural role of dopamine and the structures that receive dopaminergic innervation.

The basal ganglia. The basal ganglia comprise a number of subcortical nuclei, including the striatum. The striatum may be considered the ‘input layer’ of the basal ganglia; nearly the entire neocortex projects to it (Kemp & Powell, 1971). In turn, the striatum projects to the globus pallidus, which projects via thalamic nuclei back to the cortex; the whole makes up a ‘loop’. It is a particular characteristic of basal ganglia–thalamocortical (‘corticostriatal’) loops that although large areas of cortex send information into the loop, only a relatively small area of cortex is targeted by the return projection. Information flow in different loops is segregated — that is, the loops operate in parallel — and the loops are named for the areas of cortex to which they project: the *motor* loop (projecting in primates to the premotor cortex, supplementary motor area, and primary motor cortex, and involved in the initiation of motor acts); the *oculomotor* loop (projecting to the frontal eye fields); the *dorsolateral prefrontal* or ‘cognitive’ loop; the *lateral orbitofrontal* loop, and the anterior cingulate or *limbic* loop (projecting to the anterior cingulate cortex and medial orbitofrontal cortex) (DeLong & Georgopoulos, 1981; Alexander *et al.*, 1986). Indeed, functional segregation (parallel processing) is apparent even *within* each loop (see Alexander & Crutcher, 1990). The loops may also be differentiated on the basis of the parts of the basal ganglia and thalamus they pass through; thus, while inputs to the motor and ‘cognitive’ loops target the dorsal striatum (caudate–putamen or neostriatum), information entering the limbic loop does so through the ventral striatum. The ventral striatum consists of the nucleus accumbens (Acb), ventromedial portions of the caudate and putamen, and the olfactory tubercle; the largest component is the Acb. Within each corticostriatal loop, the basic circuitry is similar across the dorsal striatum and much of the ventral striatum (reviewed by Heimer *et al.*,

1995); it is therefore likely that the various basal ganglia loops process information in qualitatively similar ways, with the nature of the cortical target determining the apparent function of each loop.

Information processing in the basal ganglia is complex, involving not only a ‘direct’ pathway from striatum to globus pallidus (more specifically, to the internal segment of the globus pallidus and the substantia nigra pars reticulata) but a functionally antagonistic ‘indirect’ pathway from the striatum to the globus pallidus (external segment), which projects to the subthalamic nucleus, and thence to the globus pallidus (internal segment) (see Alexander & Crutcher, 1990). Cellular activity in the striatum is regulated by dopaminergic projections from the midbrain. Again, there is anatomical specificity in the dopaminergic innervation: the dorsal striatum is innervated by the SNc while the ventral striatum receives its projections from the VTA. In a further subdivision of the dorsal striatum, histochemically distinct *patches* or *striosomes* may be defined, which may project back to midbrain dopaminergic and cholinergic neurons, while the *matrix* circuitry is as described above (Grove *et al.*, 1986; Jiménez-Castellanos & Graybiel, 1989; Gerfen, 1992b; Gerfen, 1992a; Fallon & Loughlin, 1995), though it is not clear that this distinction applies to the ventral striatum (see Heimer *et al.*, 1995). In addition, there are significant dopamine projections to cortical structures that provide information to, and receive information from, the basal ganglia, such as the prefrontal cortex and amygdala (Fallon & Loughlin, 1995).

The ‘limbic loop’. This thesis will focus on the functions of the limbic loop, depicted in Figure 6. Its components include many of the structures considered part of the limbic system. The term ‘limbic’ was coined by Broca (1878) for the cortical structures encircling the upper brain stem (*limbus*, Latin for edge or border). The ‘limbic lobe’ was suggested to have a role in emotional experience and expression by Papez (1937), concepts later to be elaborated by MacLean (1949; 1952; see MacLean, 1993), who introduced the expression ‘limbic system’ to refer to the limbic lobe and its connections with the brainstem. The limbic system is not precisely defined: as the limbic lobe was considered the neural substrate for emotions, structures whose functions have to do with motivation and emotion have since been added to the anatomical definition. A modern definition of the limbic system in primates would certainly include cingulate and orbitofrontal cortex; the hippocampal formation, parahippocampal gyrus and mammillary bodies; anterior and medial thalamic nuclei; the nucleus accumbens and ventral pallidum; the amygdala and the hypothalamus.

The nucleus accumbens. On histochemical and anatomical grounds, the nucleus accumbens may be divided into the core (AcbC), shell (AcbSh), and rostral pole (a border zone with features of the other two compartments) (Zaborszky *et al.*, 1985; Zahm & Brog, 1992). The pattern of innervation of these structures differs: in terms of connectivity, Acb may be considered as having two broad functional divisions (Brog *et al.*, 1993): (1) the core, rostral pole and lateral shell; and (2) the medial shell and septal pole. Of these, the core division more closely resembles the dorsal striatum, projecting predominantly to the ventral pallidum, while the shell division also projects to subcortical structures, such as the lateral hypothalamus and periaqueductal grey, involved in the control of innate behaviours. The connections of the Acb are summarized in Table 4 and Table 5 (see Berendse *et al.*, 1992; Brog *et al.*, 1993).

As a recipient of information from a considerable array of limbic structures that projects additionally to nuclei known to be involved in behavioural expression, the Acb has been suggested to represent a ‘limbic–motor interface’ (Mogenson *et al.*, 1980). However, much of the function of the Acb is presumably related to its influence over cortical structures, the function of which are themselves somewhat mysterious. In particular, the functions of the anterior cingulate cortex (ACC) will be considered in this thesis, and the comparative anatomy of this region will be discussed in Chapter 3.

Table 4. Some inputs to the nucleus accumbens (from Brog *et al.*, 1993). Subcortical connections are nearly all reciprocal.

Region in Acb	Cortical afferents	Subcortical afferents
To all/most of the nucleus accumbens	orbital cortex posterior agranular insular cortex entorhinal cortex basal amygdala hippocampal formation (via subiculum) (Note that none of these inputs is a primary or secondary sensory area or relay.)	raphé nuclei ventral tegmental area thalamic nuclei (see Brog <i>et al.</i> , 1993 for discussion)
Shell-preferential (meaning medial shell and septal pole)	dorsal peduncular cortex infralimbic cortex pyriform cortex ventral subiculum	bed nucleus of the stria terminalis hypothalamus medial amygdala lateral habenula laterodorsal tegmental nucleus sublenticular substantia innominata lateral septal nucleus locus coeruleus
Core- or rostral pole-preferential	anterior cingulate cortex medial precentral cortex dorsal and ventral prelimbic area agranular insular cortex perirhinal cortex dorsal subiculum	dorsolateral ventral pallidum subthalamic nucleus globus pallidus substantia nigra pars compacta

Table 5. Some outputs from the nucleus accumbens (for references, see Pennartz *et al.*, 1994).

Region in Acb	Efferent connections
Core	ventral pallidum subthalamic nucleus substantia nigra pars reticulata
Shell	ventral pallidum ventral tegmental area substantia nigra pars compacta hypothalamus (preoptic, medial, lateral areas) lateral septum bed nucleus of the stria terminalis lateral habenula periaqueductal grey
Indirect, via ventral pallidum	mediodorsal thalamus pedunculopontine area (part of the mesencephalic locomotor region)

Interpretation of lesion studies

Although correlative techniques such as electrophysiology and functional neuroimaging allow the functioning of the normal brain to be measured, interventional techniques (such as lesion studies or drug infusions) are required to establish a causal link between a neural structure and an aspect of behaviour. In such studies, the anatomical specificity of the method is important. The use of aspirative or radiofrequency lesions, or local anaesthetic inactivation, will destroy or inactivate neurons in the target area, but will also affect fibres (axons) passing through the target structure, potentially affecting the function of neurons whose cell bodies are elsewhere. In the present thesis, excitotoxic lesion techniques and intracerebral drug infusions are used, both of which can affect neurons in the target site selectively. Excitotoxins typically activate NMDA-type glutamate receptors on neurons, leading to abnormal Ca^{2+} influx and cell death via apoptosis or excitotoxic necrosis; reviews have been provided by Choi (1988; 1995). Table 6 shows the conclusions that may be drawn from some of these interventional techniques.

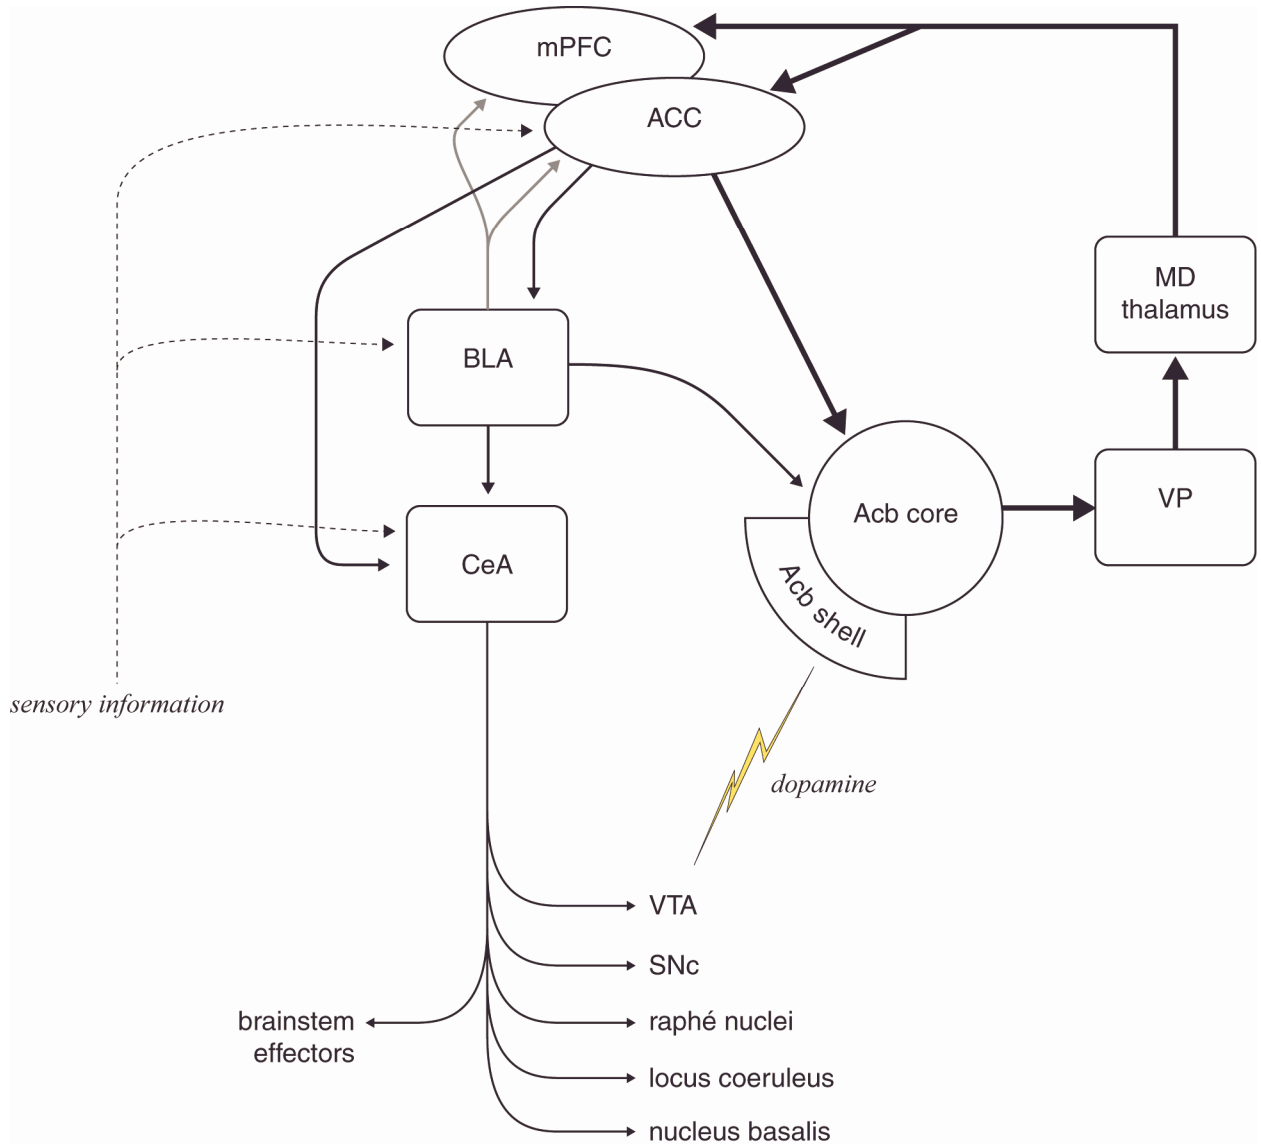


Figure 6. Part of the limbic corticostriatal loop, with associated structures. The main loop is shown in bold, together with amygdalar structures that contribute to its function in the context of appetitive approach behaviour, conditioned reinforcement and its potentiation by psychomotor stimulant drugs. For clarity, hippocampal structures are not shown. As will be discussed in the text, a functional connection between the anterior cingulate cortex (ACC) and the nucleus accumbens (Acb) core is necessary for discriminated Pavlovian approach behaviour, while the basolateral amygdala (BLA) is critical for conditioned reinforcement. The Acb is also required for the potentiation of ongoing instrumental behaviour by Pavlovian conditioned stimuli (Pavlovian–instrumental transfer; PIT). Pavlovian approach behaviour, the potentiation of conditioned reinforcement by psychostimulants, and PIT all require the central nucleus of the amygdala (CeA) and the dopaminergic innervation of the accumbens (evidence on this last point is incomplete for PIT). The integrity of the medial prefrontal or prelimbic cortex (mPFC) may be required for the perception of instrumental contingencies, and the heavy projection from the BLA to these areas of prefrontal cortex may contribute to the process by which instrumental actions are directed towards appropriate goals.

Table 6. Interpretation of lesion studies.

Manipulation	Conclusions that may be drawn from impairment	Conclusions that may be drawn from normal behaviour
Lesion, then train/test	Structure is required for learning or performance of the task	Structure is not required for learning or performance of the task, though it may still be involved
Train, lesion, test	Structure is required for performance of the task. Does not distinguish 'mnemonic' from 'motor' function.	Structure not required for performance of the task
Train in the presence of reversible inactivation; test subsequently	Either of: (a) The structure is required for task performance, and successful performance is required as part of the learning process (e.g. instrumental behaviour); (b) Learning can occur even though performance is blocked (e.g. Pavlovian conditioning), and the structure is involved in learning the task	Structure not required to learn the task
Disconnection lesion (unilateral lesion of site A and unilateral lesion of site B in the opposite hemisphere)	Site A or B must be intact bilaterally for task performance (control procedures should address this issue), or a functional connection between structures A and B is necessary for the task.	Either of: (a) A direct or indirect connection exists between the remaining A and B sites (b) Functional communication between A and B is not necessary for the task

A summary of candidate neural structures involved in Pavlovian conditioning

The focus of this thesis is firmly on the limbic structures depicted in Figure 6 (p. 35). However, to set these in a broader context, a brief review will also be provided of the role of certain other brain regions that are critical to Pavlovian conditioning.

Stimulus representation: sensory thalamus, primary and higher-order sensory cortices

The majority of sensory information concerning objects in the world reaches the brain via the sensory nuclei of the thalamus. These nuclei project directly to the amygdala, which may represent a simple, rapid route of information transfer to a structure that assesses its importance (LeDoux, 2000). But the major projections of the thalamic sensory nuclei are to primary sensory cortices, where complex attributes of stimuli are analysed. Clearly, these systems must provide CS information to Pavlovian conditioning processes; one would thus expect stimulus–stimulus associations to be organized via high-order sensory cortices. For example, cross-modal sensory preconditioning can be impaired by lesions of cross-modal sensory cortex (perirhinal cortex; Nicholson & Freeman, 2000), and following CS–food pairing, mnemonic retrieval of specific sensory aspects of the food US may depend on gustatory neocortex (see Holland, 1998). Aspects of spatial navigation, episodic memory, contextual conditioning, and other tasks requiring the integration of multiple stimuli into a 'scene' are sensitive to hippocampal lesions (e.g. Gaffan & Harrison, 1989; Selden *et al.*, 1991; Gaffan, 1992; Phillips & LeDoux, 1992; Maren, 1999). Pavlovian conditioning can also lead to increases in the cortical representation of significant CSs (Weinberger, 1995; 1998a; 1998b).

The amygdala

Review

The amygdala has long been implicated in 'affective' Pavlovian conditioning, in both appetitive and aversive settings (Davis, 1992; Everitt & Robbins, 1992; LeDoux, 1992; Holland, 1997; Everitt *et al.*, 2000a; LeDoux, 2000). Two of its major components are the basolateral nuclear group (BLA) (including the basal and lateral nuclei) and the central nucleus (CeA) (for anatomical reviews, see Amaral *et al.*, 1992; Pitkänen, 2000).

The BLA is heavily implicated in Pavlovian fear conditioning. In a typical fear conditioning experiment, an auditory or visual CS is paired with electric shock to the feet; a CR rapidly develops in which the rat freezes for the duration of the CS (a species-specific defence reaction), while a number of autonomic changes occur. Lesions of the BLA dramatically impair the conditioned freezing response. The BLA receives information about CSs from sensory thalamic nuclei and cortices, and information about the painful US from somatosensory thalamus and cortex, while synaptic plasticity has been demonstrated in the BLA during conditioning and is necessary for conditioning to occur (for review, see LeDoux, 2000). In turn, the BLA sends a heavy projection to the CeA, which projects to a wide array of hypothalamic and brainstem structures, including the chemically-defined projection systems of the isodendritic core such as the noradrenergic locus coeruleus, the dopaminergic SNc and VTA, the serotonergic raphe nuclei and the cholinergic nucleus basalis magnocellularis (of Meynert; NBM) (Amaral *et al.*, 1992; Davis, 1992). Through these and other targets, the CeA can command autonomic responses, trigger simple skeletomotor acts such as the rat's freezing response, and regulate attentional and arousal function. It is undoubtedly the case that the BLA, which does not project directly to these structures, uses the CeA for this purpose. Thus, the dominant model of conditioning in the amygdala (Davis, 1992; LeDoux, 1992; LeDoux, 2000) suggests that CS–US association occurs in the BLA, which then expresses CRs through the CeA.

However, this picture is incomplete. Firstly, the BLA has independent projections to the ventral striatum and prefrontal (particularly orbitofrontal) cortex, giving it access to more complex response mechanisms (Everitt & Robbins, 1992; Everitt *et al.*, 1999). Secondly, the CeA also receives polymodal sensory information from association cortex (McDonald, 1998), giving it the potential to form CS–US associations, and a number of studies have demonstrated double dissociations between the effects of BLA and CeA lesions on measures of conditioning, both appetitive and aversive.

The first such demonstration concerned the modulation of instrumental behaviour. If rats are trained to respond on two concurrent schedules of food reinforcement, and one schedule additionally produces a CS terminating in mild footshock, they learn to bias their responding away from the punished lever, but also exhibit conditioned suppression when the CS occurs. Killcross *et al.* (1997b) demonstrated that BLA lesions impaired rats' ability to direct their instrumental behaviour away from the punished lever, leaving conditioned suppression intact, while CeA lesions had exactly the opposite effect (preserved instrumental avoidance with abolished conditioned suppression). A similar double dissociation using an appetitive version of the task was recently reported (Killcross *et al.*, 1998). From demonstrations such as these, a new view of amygdala function has emerged; the data are reviewed in full by Everitt *et al.* (2000a) and a summary is presented below.

The basolateral amygdala: retrieval of the current value of the US

A great deal of evidence has accumulated showing that rats with BLA lesions can acquire first-order conditioned responses, but that these responses are insensitive to reinforcer revaluation. For example, rats with BLA lesions have been shown to acquire normal conditioned responding to a CS paired with food (the CR being approach to the cup into which food was delivered). BLA-lesioned rats also showed normal acquisition of an aversion to that food when it was subsequently paired with LiCl, but failed spontaneously to adjust their responding (orienting and food cup approach) to the CS after the food was devalued (Hatfield *et al.*, 1996). Similar results have been observed in monkeys (Málková *et al.*, 1997). The most parsimonious explanation is that the conditioned responses learned by the BLA-lesioned rats were a result of direct associations between the CS and the response. They lacked the ability to use the CS to access the

value of a specific US and use that representation to alter their response. Holland (1998) defines this ability as ‘mediated performance’: the ability to respond based on a CS-activated representation of the US.

The idea that BLA-lesioned animals cannot use a CS to gain access to the current value of its specific US has great explanatory power. Second-order conditioning requires that the second-order stimulus becomes associated with the affective value that is called up by the first-order CS (as discussed earlier, p. 21) (see also Gewirtz & Davis, 1998): BLA-lesioned rats cannot acquire second-order conditioning (Hatfield *et al.*, 1996), cannot acquire responding under second-order instrumental schedules (Everitt *et al.*, 1989; Whitelaw *et al.*, 1996), and cannot use a first-order CS as a conditioned reinforcer (Cador *et al.*, 1989; Burns *et al.*, 1993). Clearly, the responses that still occur to the first-order CS do not support second-order conditioning, while the effects on reward devaluation (Hatfield *et al.*, 1996) demonstrate that the deficit in BLA-lesioned animals is not restricted to second-order conditioning *per se*. Specific modulation of instrumental choice behaviour by a CS also requires that the subject utilizes the motivational value of a particular US; this capability, too, depends upon the BLA (Killcross *et al.*, 1997b; Killcross *et al.*, 1998).

The formation of an association between a CS and the affective value of a US also accounts for responses such as conditioned freezing, which cannot readily be accounted for in terms of a CS–UR association. Thus, the conditioned freezing response does not resemble the UR to shock, which is characterized by agitation, jumping, vocalization and escape, but instead represents an adaptive response to danger. At the time of conditioning, therefore, there is no freezing response occurring to which a CS–UR association can be formed (see Wagner, 1970, p. 154, for discussion of this). In addition, freezing is a US-specific conditioned response: while freezing occurs to a CS for shock, it does not occur to a CS for the omission of expected food, even though both signal aversive events (as discussed above, p. 21). It seems plausible to suggest, therefore, that the BLA is critical for the acquisition of conditioned freezing because it subserves the formation of an association between the CS and a neural representation of the affective properties of the US (Bolles & Fanselow, 1980). Similarly, fear-potentiated startle may reflect the potentiation of a reflexive startle response by an affective representation retrieved by the CS, and is thereby sensitive to BLA lesions (Davis, 1997; Walker & Davis, 1997).

The central nucleus of the amygdala: stimulus–response associator and controller of the brainstem

Even though it receives neuronal afferents appropriate to support them, there is no direct evidence to suggest that the CeA is itself a site of CS–US associations; it might receive an already-associated input. However, it is clear that animals lacking a BLA can form some kinds of association, the conditioned expression of which is sensitive to CeA, but not BLA, lesions (Gallagher & Holland, 1994; Killcross *et al.*, 1997b; Hall *et al.*, 1999; Parkinson *et al.*, 2000b). The simplest analysis at present seems to be that the CeA does form simple CS–UR (‘sensorimotor’) associations, which do not depend upon a specific US: that is, they are independent of the identity and current motivational value of the US and are also unable to support second-order conditioning. We have suggested (Everitt *et al.*, 2000a) that the responses subserved by CeA-dependent associations especially include the modulation of reflexes organized within the brainstem, including some that might conventionally be regarded as ‘affective’, including conditioned suppression, conditioned orienting, and Pavlovian–instrumental transfer. These are all disrupted by CeA but not BLA lesions. Responses such as conditioned suppression may influence instrumental behaviour non-specifically, but are insufficient to modulate instrumental behaviour differentially, as assessed in choice tasks (Killcross *et al.*, 1997b).

Additionally, Gallagher, Holland and co-workers have also shown that through its projections to the reticular formation, the CeA is involved in the control of attentional aspects of stimulus processing. The

CeA plays a role in visuospatial attention during continuous-performance tasks (Holland *et al.*, 2000), and also appears to regulate the *associability* of stimuli under certain circumstances (Gallagher & Holland, 1992; Gallagher & Holland, 1994; Holland & Gallagher, 1999). Associability is a learning-theory concept (e.g. Rescorla & Wagner, 1972; Pearce & Hall, 1980); it determines how much processing is devoted to a CS, and therefore indirectly determines the degree to which new things can be learned about the CS. The Pearce & Hall (1980) model of Pavlovian conditioning suggests that when a CS is reliably followed by a US, the CS may be worth responding to, but is not worth learning about: animals should confine their attention to learning about stimuli whose consequences are less well known. Associability can be increased by surprising events: for example, if a light is regularly followed by a tone, presentation of the light on its own (with the surprising absence of the tone) is predicted by the Pearce–Hall model to increase the subsequent associability of the light (e.g. Wilson *et al.*, 1992; see Holland, 1997). This phenomenon — specifically, the ability to *upregulate* associability — appears to depend upon the integrity of the CeA (Holland & Gallagher, 1993b; 1993a), together with its projections to cholinergic neurons in the NBM and substantia innominata (Han *et al.*, 1999), and possibly from there to the posterior parietal cortex (see Holland, 1997). Though the cellular basis of associability is unknown, it is interesting to note that Weinberger and colleagues have shown that auditory cortex receptive fields for a CS of a particular frequency expand, at the expense of other regions, when that CS is paired with an aversive US. This cortical plasticity depends upon muscarinic acetylcholine receptors and can be induced by stimulation of the NBM (see Weinberger, 1995; 1998a; 1998b). Expansion of a sensory receptive field might be one mechanism by which the associability of a stimulus could increase, as might increased attention to that stimulus directed by the attentional circuits known to exist in the posterior parietal cortex (see Posner, 1995).

Summary

This view of amygdala function is illustrated speculatively in Figure 7. When a CS predicts an appetitive US, it may form associations with sensory and motivational representations of that US (links 1 and 2 in the figure), with central affective states (3) and with unconditioned responses at some level (4). When the US is devalued, its motivational representation is in some way selectively redirected to an aversive state (not shown), so it is through link 1 or 2 that the changed response to a first-order CS occurs. It should be noted that while affective states are illustrated as ‘centres’, very little is known of the neuronal mechanism by which valence might be encoded: such information might just as easily be carried as a temporal or chemical code and be multiply represented, rather than existing in distinct spatial loci. Indeed, it has been convincingly argued that the orbitofrontal cortex (OFC), which has extensive reciprocal connections with the BLA and has also been implicated in CS retrieval of US value (Gallagher *et al.*, 1999), provides an important site for the representation of affective valence (Schoenbaum *et al.*, 1998; Rolls, 1999; Schoenbaum *et al.*, 1999; Rolls, 2000). The exact relationship between BLA and orbitofrontal function is not clear at present; further data are reviewed when considering instrumental behaviour, below.

It is at present unclear whether the BLA is involved in representing specific sensory information about USs, required for S–S associations. Each sensory modality projects to a region of sensory cortex, a reason to question a priori whether the BLA is required, and rats can learn stimulus discrimination tasks in the absence of the BLA (Schwartzbaum, 1965; Sarter & Markowitsch, 1985; Burns *et al.*, 1999). If the BLA is involved, it would therefore have to be as an ‘independent associator’ (E in Figure 2, p. 21). According to this scenario, BLA-lesioned animals make unconditioned responses and learn simple CS–UR associations, including ‘emotional’ responses, but the CS would convey no information about the identity of the US. Alternatively, the US-specific representation involving the BLA might be purely affective; in this alternative scenario, BLA-lesioned animals can learn CS–UR associations and CS–US(sensory) associa-

tions, but cannot learn CS–US(affective) associations, and the sensory representation they can activate is without affective valence (see also Holland, 1998, for a discussion of this possible dissociation). Following a recent demonstration that BLA lesions do not impair sensory preconditioning (Blundell & Killcross, 2000b), the latter interpretation seems most likely.

It is also presently unclear whether the BLA holds US-specific representations that excite general appetitive/aversive states in another structure, or itself contains this ‘affective processor’, or contains both. It is clearly difficult to distinguish whether BLA-lesioned animals lack affective states that may take part in associations, or merely cannot call them up via a CS; however, transreinforcer blocking and performance (but not acquisition) of second-order conditioning are two phenomena that appear to depend on simple affect, so further experiments may resolve these issues.

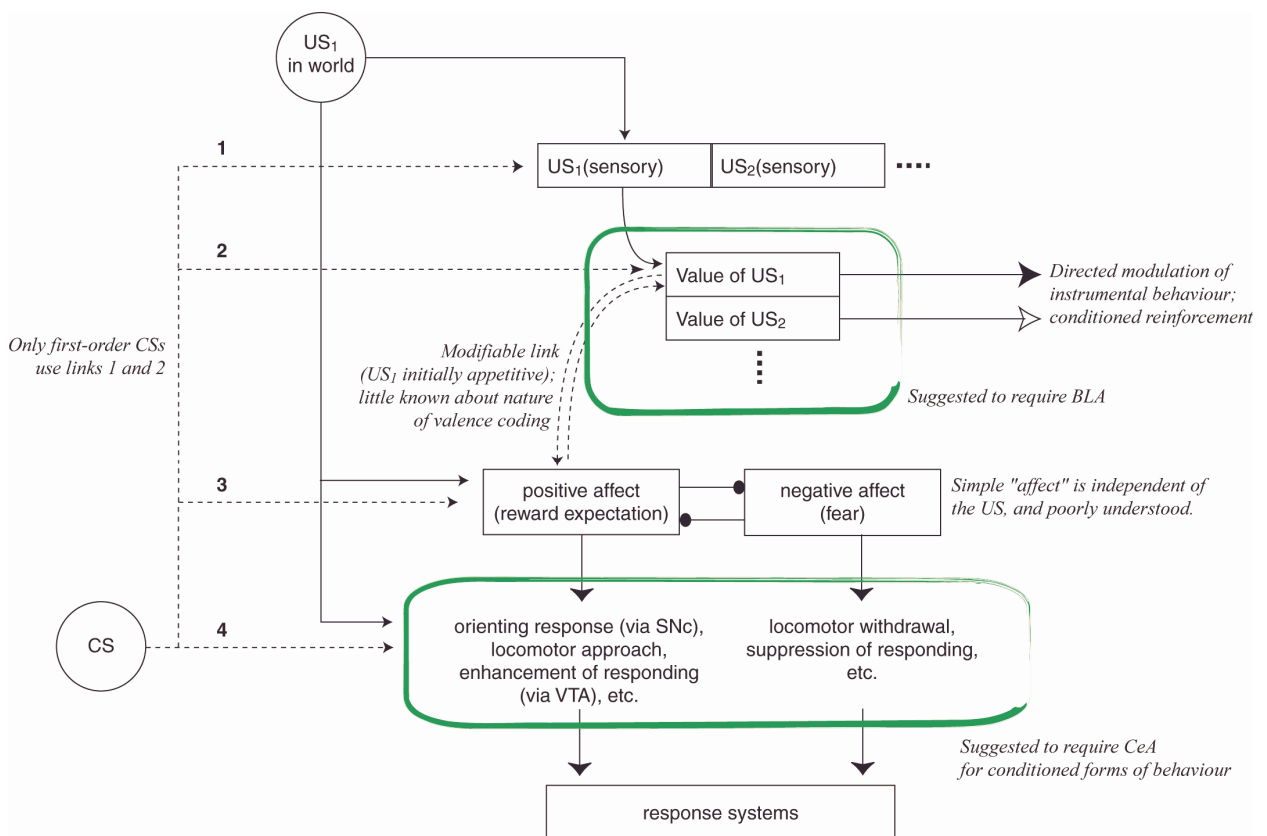


Figure 7. Schematic of representations that may be involved in Pavlovian conditioning, emphasizing the hypothesized role of amygdaloid subregions. The BLA is required for a CS to gain access to the current value of its specific US. In the figure, the CS has been associated with US₁, initially appetitive, while an unrelated US₂ maintains a separate value (connections not shown for clarity). As discussed in the text, the precise nature of the information encoded in the BLA is uncertain; here, it is illustrated as binding US-specific sensory information to an affective value. The BLA may use this information to control CeA function but also to modulate specific instrumental (choice) behaviour, as in conditioned reinforcement tasks; the Acb is a key target of this information. In contrast, the CeA is required for CS–UR learning, particularly when the response involves modulation of hypothalamic and brainstem functions. The CeA may also modulate the associability of CSs (see text), but this function is not illustrated.

Conditioning of simple CS–UR skeletal responses with high temporal precision: the cerebellum

It would be elegant if the representations encoded by amygdalar nuclei could be entirely categorized using a well-defined psychological dichotomy. It appears that we are remarkably close to this situation with the suggestion that the CeA encodes or expresses Pavlovian stimulus–response (CS–UR) associations, while the BLA encodes or retrieves the affective value of the predicted US. However, not all stimulus–response associations depend on the CeA. For example, nictitating membrane/eyeblink conditioning de-

ponse associations depend on the CeA. For example, nictitating membrane/eyeblink conditioning depends instead on the cerebellum, even though the eyeblink clearly is part of the UR to eyeshock; this circuit has been extensively mapped (see e.g. Thompson *et al.*, 2000) and appears to involve CS–UR associations. Eyeblink conditioning can occur in the absence of the amygdala (even though simultaneously conditioned changes in heart rate are amygdala-dependent). In attempting to define the purview of cerebellar conditioning, Steinmetz (2000) comes to a more pragmatic, neurobiological solution: the cerebellum has been shown to be involved in associative learning when (1) a simple motor response is involved; (2) the CS–US interval is shorter than ~4 seconds; (3) the US is aversive; (4) the US not only causes a UR, but *in addition* activates the inferior olive, the ‘teaching system’ for such cerebellar learning. This definition fits no neat psychological category so far proposed. Applying this rationale to the amygdala, for example, would lead to the suggestion that the CeA subserves Pavlovian CS–UR associations when that response is controlled by a hypothalamic or brainstem nucleus governed by the CeA; such responses include autonomic changes, motivational arousal and attentional enhancement.

This observation has implications for general theories of learning. Belief in a general learning process has justification (Dickinson, 1980, pp. 6–9), and has led to undoubted success in describing conditioning phenomena. If associations in the cerebellum are formed according to different rules to associations in the CeA, however, there is no universal learning process. On the other hand, if such disparate systems *do* learn according to the same rules of association, why? This would imply either that highly complex associative rules are embedded on a small scale (such as at the level of the neuron) in a wide variety of neural tissue, and very consistently so, or that some as yet unknown central, cooperative learning mechanism regulates learning in widely distributed areas of the brain. There is direct psychological evidence for the latter idea (see Wagner, 1978; Dickinson, 1980, chapter 4; Baars, 1988), and the elucidation of the neural basis of this mechanism is an exciting challenge.

The anterior cingulate cortex and stimulus-reinforcer associations

One other structure within the ‘limbic loop’ of Figure 6 (p. 35) that has been implicated in Pavlovian conditioning is the ACC. The ACC was first implicated in aversive conditioning: it receives excitatory nociceptive information from midline and intralaminar thalamic nuclei (Hsu & Shyu, 1997), and is capable of commanding autonomic responses (Fisk & Wyss, 1997). Early studies found that aspirative lesions of the ACC, which also destroy fibres of passage, attenuated classically conditioned bradycardia in rabbits (Buchanan & Powell, 1982a). More recently, the role of the rat ACC in appetitive Pavlovian conditioning has been studied, using excitotoxic lesions (to prevent damage to fibres of passage) and the phenomenon of autoshaping.

Autoshaping. In the autoshaping task that has been studied in our laboratory, a visual stimulus (CS+) is presented on a computer screen and is followed by the delivery of food in a *different* spatial location. A second stimulus (CS–) is also presented, but never followed by food. Though the subject’s behaviour has no effect on food delivery, animals develop a conditioned response in which they selectively approach the CS predictive of food before returning to the food hopper to retrieve the primary reward. Autoshaping is generally considered to be a Pavlovian conditioned response, as it can be acquired even under an omission contingency in which approach to the CS+ prevents food delivery (Williams & Williams, 1969) (see also Davey *et al.*, 1981). There is an alternative possibility: that the response is instrumental and is shaped by contact with the CS, which acts as a conditioned reinforcer (see Williams, 1994a, p. 471); nevertheless, in either situation, the Pavlovian relationship between the CS and primary reward underlies autoshaping.

Using this task, Bussey *et al.* (1997a) found that ACC-lesioned rats exhibited a profound impairment in the acquisition of autoshaped responding. Lesioned rats were also impaired on an appetitive discrimination task requiring rats to learn eight stimulus–reward associations concurrently (Bussey *et al.*, 1997b). Bussey *et al.* (1996) also found that ACC lesions *facilitated* early learning of a conditional visual discrimination (CVD) task that was soluble by the formation of stimulus–response habits, but not by the formation of stimulus–reinforcer associations; Bussey *et al.* suggested that the formation of stimulus–reinforcer associations hinders performance on this task by competing with a stimulus–response system. Taken together, these studies suggest strongly that the ACC is involved in some aspect of Pavlovian stimulus–reinforcer association.

Expression of Pavlovian conditioning: amygdala-accumbens and cingulate-accumbens interactions

Though the ACC is required for the development of autoshaping, this task also requires the AcbC (though not the AcbSh) (Parkinson *et al.*, 1996; Parkinson, 1998). Furthermore, AcbC lesions impair the performance of the conditioned response in rats trained before the lesion was made (Everitt *et al.*, 2000b), just as they impair temporally discriminated Pavlovian approach to a single CS predictive of food (Parkinson *et al.*, 1999b). Similarly, 6-OHDA-induced dopamine depletion of the whole Acb impair both the acquisition (Parkinson *et al.*, submitted) and performance (Everitt *et al.*, 2000b) of autoshaping.

Indeed, the ACC is the only major limbic cortical afferent to the Acb that is required for autoshaping, as lesions of BLA, dorsal or ventral subiculum, medial prefrontal cortex (mPFC), or posterior cingulate cortex (PCC) have no effect on its acquisition (Parkinson *et al.*, 1996; Bussey *et al.*, 1997a; Parkinson *et al.*, 2000b).

It seems likely, then, that stimulus–reward associations stored or retrieved by the ACC gain behavioural expression through the Acb. This hypothesis was tested directly by Parkinson *et al.* (2000c), who used a ‘disconnection’ procedure, in which asymmetric unilateral lesions of both the ACC and Acb were made in order to prevent communication between the two structures. The disconnection impaired autoshaping, though single unilateral lesions of either structure did not.

Autoshaping is not the only form of Pavlovian conditioning in which the Acb appears to give behavioural expression to associative information arising from limbic cortical or quasi-cortical afferents. At least three other tasks have been shown to operate similarly. The first is the expression of a conditioned place preference; this depends on the BLA, but also on the Acb, and a disconnection lesion of the two structures abolishes behavioural expression (Everitt *et al.*, 1991). The second is second-order conditioned approach: Setlow *et al.* (2000) recently demonstrated that BLA–Acb disconnection impairs the acquisition of second-order conditioned approach behaviour, but not second-order conditioned orienting, or first-order conditioned approach — consistent with the known involvement of the BLA in second-order conditioning (Hatfield *et al.*, 1996), and the Acb in conditioned approach (Parkinson *et al.*, 1999b; Everitt *et al.*, 2000b; Parkinson *et al.*, 2000c). The third is responding for conditioned reinforcement, discussed further below. Briefly, lesions of the BLA impair responding for conditioned reinforcement (Burns *et al.*, 1993); injection of amphetamine into the Acb dramatically enhances responding for conditioned reinforcement (Taylor & Robbins, 1984; Burns *et al.*, 1993), and this enhancement depends on the integrity of the BLA — again suggesting expression of amygdala-dependent information via the Acb.

Projections from the ACC and BLA to the Acb are direct and glutamatergic. In addition to these highly specific, information-rich projections, recent evidence suggests that the amygdala — specifically, the CeA — may also modulate Acb function via a different route. The CeA does not project to the AcbC (Zahm & Brog, 1992; Brog *et al.*, 1993; Parkinson, 1998) or the AcbSh (Zahm *et al.*, 1999, pp. 1119/1124), but does project to the VTA, the source of the dopaminergic innervation of the Acb (Amaral

et al., 1992, p. 35; Fudge & Haber, 2000). It may be that the CeA is capable of regulating Acb DA; in accordance with this hypothesis, lesions of the CeA impair the acquisition of autoshaping (Parkinson *et al.*, 2000b) and the potentiation of responding for conditioned reinforcement by intra-accumbens amphetamine (discussed further below; Robledo *et al.*, 1996), tasks that depend on Acb DA (Taylor & Robbins, 1986; Parkinson *et al.*, submitted).

The anterior cingulate cortex: unanswered questions

Evidence has been summarized that the ACC is involved in Pavlovian conditioning. However, this evidence stems from a limited range of tasks, and it is unclear exactly what role the ACC plays. The pattern of impairment that Bussey *et al.* (1997a) found in autoshaping was not complete absence of responding to the CS+ or the CS–, but rather a loss of discrimination through increased CS– responding. Given the known role of the prefrontal cortex in response inhibition (summarized by Roberts *et al.*, 1998), it is possible that the ACC plays a unique inhibitory role within the ‘limbic loop’ in the expression of Pavlovian conditioning. Similarly, it is not known whether the ACC is critical for the expression of autoshaping, as well as its acquisition. Moreover, there are a number of other gaps in the story. It is not known whether the ACC is involved in all forms of conditioned approach, whether (given its relationship to the BLA and the Acb) it is another important contributor to the effects of conditioned reinforcers or other Pavlovian influences on instrumental performance, or whether excitotoxic lesions of the ACC impair aversive as well as appetitive conditioning. These questions will be addressed in Chapter 3, where a full review will also be made of the comparative anatomy and functions of the ACC across different species.

A summary of candidate neural structures that influence instrumental performance

Instrumental performance: outside the ventral striatum

The multifactorial view of instrumental performance outlined earlier is relatively young (see Dickinson, 1994; Dickinson & Balleine, 1994); as a result, few studies have as yet investigated the neural basis of theoretically well-defined processes contributing to instrumental performance. Of necessity, then, this summary will be brief, and will focus on the contribution of the Acb.

Contingency detection: medial prefrontal cortex and dorsal hippocampus

Demonstration that a structure is necessary for detection of action–outcome contingencies requires more than showing that an animal cannot acquire instrumental responding in its absence. Indeed, were one to prevent an animal from perceiving contingencies, there is every reason to think that instrumental performance *would* be acquired, via a habit system. Explicit tests of contingency perception are thus required. For example, rats may be trained to perform two actions concurrently for two different food rewards; in addition, one of those reinforcers may be delivered noncontingently with respect to the subjects’ behaviour. The degree of action–outcome contingency for this reinforcer, $P(\text{outcome} \mid \text{action}) - P(\text{outcome} \mid \text{no action})$, is thus selectively degraded. In one of the few lesion studies to use this technique, Balleine & Dickinson (1998a) found that although lesions of prelimbic cortex did not prevent rats acquiring instrumental performance, or, in separate tests, from discriminating between the two actions and the two reinforcers, they rendered the rats insensitive to this contingency manipulation, suggesting that such rats might truly be ‘creatures of habit’. Similar results have been obtained with lesions (albeit electrolytic) of the dorsal hippocampus (Corbit & Balleine, 2000b). However, hippocampal lesions appeared not to impair contingency learning *per se* but instead impaired sensitivity to the background, noncontin-

gent reinforcement; these effects may have resulted from a failure of contextual conditioning (Corbit & Balleine, 2000b) (see p. 36).

Indirectly, these experiments also provide an insight into habit formation, for they imply that a habit can develop without the subject ever possessing knowledge about the instrumental contingency. These studies suggest that habits 'build up' independently of goal-directed action, likely as a function of the number of reinforcers received (Adams, 1982), and are *exposed* in normal animals at the point that the action–outcome contingency no longer controls performance (Dickinson, 1985), though studies of the dysfunctioning brain cannot rule out the possibility that interaction between the two systems occurs when the intact animal learns.

Incentive value: insular cortex and other candidates — hypothalamus, ventral pallidum, amygdala, orbitofrontal cortex

Balleine & Dickinson (1998a; 2000) also investigated the role of the insular cortex, the primary gustatory cortex in the rat (Norgren, 1995), in incentive learning for food rewards. Lesioned rats performed normally on the contingency test described above. In addition, a specific satiety test was conducted, in which the rats were fed one of the foods to satiety, thus giving them the opportunity to learn that this food had reduced value in the sated state. The rats only ever experienced the other food whilst hungry. Finally, the rats' instrumental performance was tested in extinction while sated. While sham-operated control rats responded less for the reward that had been devalued, insula-lesioned rats failed to make this discrimination. However, in a further test in which the reinforcers were actually delivered, they discriminated immediately. This suggests that the insula cortex is not a critical structure for determining incentive value *per se*, but is critical for storing or retrieving the memory of the incentive value in the absence of the reward. As incentive value can be retrieved via tastes (Rescorla, 1990b; Balleine & Dickinson, 1998b), this accords with the known functions of insular cortex (see p. 36 for a similar view derived from studies of Pavlovian conditioning), although it implies some degree of dissociation between primary perception of taste and taste memory.

Incentive learning depends upon the availability of information regarding the motivational state of the animal. Are there any obvious neural candidate providers of such information? The hypothalamus, in the ventral forebrain, is such a candidate. The hypothalamus serves as the final controller of diverse bodily homeostatic systems, including endocrine function (via the pituitary gland), thermoregulation, autonomic control, and circadian rhythmicity. It plays a key role in initiating 'consummatory' behaviours, such as eating, copulation, and acts of aggression (for reviews see Swanson, 1987; Simerly, 1995). It is also the brain region responsible for detecting many of the physiological variables relevant to motivational states such as hunger and thirst; for example, it responds to blood glucose levels, gut hormones released in response to feeding, tissue osmolality, and systemic hormones released in response to fluid depletion (reviewed briefly by Kupferman, 1991). Indeed, the gut 'satiety hormone' cholecystokinin (CCK) has been shown to affect incentive learning directly (Balleine & Dickinson, 1994; Balleine *et al.*, 1995a), as benzodiazepines do (Balleine *et al.*, 1994).

If the incentive learning hypothesis presented above is correct, the next stage in the assignment of incentive value is hedonic experience. As discussed earlier, there are only limited techniques available for assessing the hedonic impact of foods. If the taste reactivity test (Grill & Berridge, 1985) is accepted as a measure of hedonic impact, a variety of anatomical and neurochemical systems contribute to hedonic experience, including opioid systems in the AcbSh, benzodiazepine-sensitive systems in the brainstem, and a possible common pathway in the ventral pallidum (reviewed by Berridge, 1996; Berridge & Robinson,

1998, pp. 316–317), though dopaminergic systems appear not to play a role in hedonic experience (Berridge & Robinson, 1998).

Two other major structures have been implicated in the representation of value and the control of behaviour. These are the amygdala and orbitofrontal cortex.

As discussed above, both the BLA and CeA contribute to ‘affective’ Pavlovian responses; the BLA is suggested to be critical for a process by which a CS retrieves the affective value of a US, and for directing instrumental behaviour accordingly (Everitt *et al.*, 2000a). A prime example of such direction is responding for conditioned reinforcement, in which the BLA directs the selection of actions according to the acquired value of the conditioned reinforcer (Cador *et al.*, 1989; Burns *et al.*, 1993). Although conditioned reinforcers may have multiple attributes, it is at least reasonable to suggest that they possess incentive value (see p. 31).

Tests of contingency perception and incentive learning in BLA-lesioned rats (A. Dickinson and B.W. Balleine, unpublished observations) are consistent with the hypothesis that the BLA is involved in determining incentive value, but are not conclusive. In the specific satiety test described above, BLA-lesioned rats were impaired at discriminating between the devalued and valued reinforcers, both in an extinction test and a reinforced test (see also Málková *et al.*, 1997). However, they were also insensitive to the contingency test described above, suggesting either (1) that they could not perceive the contingency, (2) that they could not discriminate the actions, or (3) that they could not discriminate the reinforcers. BLA-lesioned rats performed normally in a different form of contingency test, using a single reinforcer, in which subjects had to perform two actions (A1 and A2) in the specific order A1→A2 in order to obtain food (a task described in Balleine *et al.*, 1995b). Just like sham-operated rats, BLA-lesioned rats selectively increased the probability of making the chained response A1→A2 compared to the three other possible response patterns (A1→A1, A2→A2, A2→A1). This suggests that they could at least discriminate the two actions, but it does not rule out deficits in reinforcer discrimination, making the BLA’s specific contribution to incentive value unclear.

In this respect, the other connections of the BLA should not be neglected. For example, the connection between the BLA and the mPFC has recently been shown to be involved in the ability of rats to modulate instrumental choice behaviour in response to conditioned punishment (Coutureau *et al.*, 2000); the anatomical connection between the BLA and the mPFC (Pitkänen, 2000) might conceivably represent a functional link between incentive value and instrumental contingencies. Additionally, the BLA is extensively and reciprocally connected to the orbitofrontal cortex (reviewed by Öngür & Price, 2000), which has been widely suggested to guide behaviour based on the anticipated value of different actions (Nauta, 1971; Damasio, 1994). In primate orbitofrontal cortex, cells may be found that respond to reward but discriminate between different rewards in doing so (Schultz *et al.*, 1998; 2000). The orbitofrontal cortex is a particularly strong candidate for a representation of incentive value, as its neurons respond rapidly to changes in the reward value of specific foods. For example, when a monkey is fed to satiety with a particular food, the orbitofrontal cortex responds to its flavour or odour decline, while the responses to other foods are unaffected (see Rolls, 2000), paralleling the behavioural change induced by sensory-specific satiety. Like the amygdala, the orbitofrontal cortex is well placed to process specific value information, as it receives projections from polymodal sensory cortex in addition to motivational state information from the hypothalamus. The relationship between the orbitofrontal cortex and the amygdala is at present unclear. Although Rolls has suggested that primate orbitofrontal cortex acts as a highly flexible system that takes over functions of the more primitive amygdala (Rolls, 2000), Schoenbaum *et al.* (1999) found evidence that, in the rat, the BLA rapidly learns to respond to CSs according to the motivational value of the

US, while changes in the electrophysiological response of orbitofrontal cortex cells follow later and are more clearly related to choice behaviour. Recently, direct evidence for a functional connection between the BLA and orbitofrontal cortex has been provided by Baxter *et al.* (2000), who showed that disconnecting these two structures impaired the ability of rhesus monkeys to adjust their choice behaviour in response to reinforcer devaluation.

Stimulus–response habits: the dorsal striatum?

The dorsal striatum (a component of the ‘motor loop’ of the basal ganglia; Alexander *et al.*, 1986; Alexander & Crutcher, 1990; Alexander *et al.*, 1990), together with its dopaminergic innervation, facilitates stimulus–response coupling — that is, the initiation of motor actions in response to environmental stimuli (Robbins & Everitt, 1992). It is natural to suggest that such stimulus–response coupling may underlie S–R habits in instrumental behaviour. A role for the basal ganglia in habit formation was originally suggested by Mishkin *et al.* (1984), who saw a habit as a direct stimulus–response association that was learned slowly but was stable. A recent review of this concept is provided by White (1997). Much of the subsequent work on this issue has proved controversial (see Wise, 1996; Wise *et al.*, 1996), and some of the best evidence for a long term change in behaviour that is dependent on the striatum is from an experiment by Packard & McGaugh (1996). They trained rats in a T-maze with one arm consistently baited. This task is soluble by two methods: repeating the reinforced response (the physical response of turning left or turning right), or approaching the place where food was found (a ‘place response’). These may be distinguished by letting the rat approach the choice point from the opposite direction. After 8 days of training, most rats made place responses. Inactivation of the dorsal hippocampus with lidocaine (lignocaine) on the test session eliminated this tendency, such that the rats showed neither place nor motor response learning, but inactivation of the dorsolateral caudate nucleus had no effect. After 16 days of training, however, most rats made the motor response that had been reinforced. Inactivation of the hippocampus had no effect, whilst inactivating the caudate eliminated ‘motor’ responding and reinstated place responding. Therefore, development of a stimulus to motor response mapping takes place slowly during reinforced training and comes to dominate behaviour, and its performance depends on the caudate nucleus. However, it should be noted that studies like this one do not always satisfy the definition of ‘habit’ used in the discussion of instrumental behaviour above. For example, few such studies have tested the effect of reinforcer devaluation on performance of the presumptive habit.

The nucleus accumbens, dopamine, and the impact of Pavlovian conditioned stimuli

Goal-directed action does not require the nucleus accumbens

The available evidence suggests that the Acb is not required for goal-directed action. Balleine & Killcross (1994) studied rats with excitotoxic lesions of the Acb performing a lever-pressing task. They established that these rats remained sensitive to a change in the instrumental contingency (from response-contingent to non-contingent reinforcer delivery); in addition, they were sensitive to a change in the value of the instrumental outcome. By the criteria of Dickinson & Balleine (1994), these rats remained capable of goal-directed action. Similarly, dopamine receptor antagonists do not affect the representation of value that governs goal-directed action (the instrumental incentive value; Dickinson *et al.*, 2000). Insofar as the issue has been addressed experimentally, stimulus–response habits also persist following Acb lesions or dopamine depletion (Robbins *et al.*, 1990a; Reading *et al.*, 1991), although these studies did not use outcome devaluation tests to demonstrate that behaviour was habitual.

At first sight, these results are inconsistent with studies showing that manipulations of Acb affect responding for food. For example, Kelley *et al.* (1997) demonstrated that NMDA receptor blockade of the AcbC impaired the acquisition of a lever-press response for food, though not its subsequent performance on an VR2 schedule. Similarly, Salamone and colleagues have shown that dopamine depletion of Acb reduces the ability of to perform instrumental responses when the work requirement is high (e.g. Aberman & Salamone, 1999). Indeed, Balleine & Killcross (1994) found that Acb-lesioned rats responded at a lower level than controls.

However, when simple reinforcement schedules are used, there are many potential influences on performance. One such influence is the impact of Pavlovian conditioned stimuli (CSs) in the environment, and, as suggested by Balleine & Killcross (1994), it is for the impact of these stimuli that the Acb appears critical. As discussed above (p. 42), the Acb is required for autoshaping, in which locomotor behaviour is controlled by appetitive Pavlovian conditioned stimuli. In addition, it has been shown to be involved in at least two situations in which Pavlovian stimuli affect instrumental behaviour.

Responding for conditioned reinforcement is affected by accumbens manipulations

Following the suggestion by Hill (1970) that an important mechanism of action of psychostimulant drugs was to enhance the effects of conditioned or secondary reinforcers, amphetamine was shown to potentiate responding for conditioned reinforcement when injected directly into the Acb (Taylor & Robbins, 1984). In the prototypical task, rats are first trained to associate a CS with the delivery of primary reinforcement. In a subsequent extinction test, they are presented with two levers; responding on the CRf lever results in delivery of the CS, while responding on the NCRf lever has no consequence. Intra-accumbens dopamine agonists greatly enhance responding for the conditioned reinforcer, an effect that is anatomically, behaviourally and pharmacologically specific (Taylor & Robbins, 1984; Taylor & Robbins, 1986; Cador *et al.*, 1991).

Subsequent studies have demonstrated that the ability of amphetamine to potentiate responding for conditioned reinforcement depends on the integrity of the AcbSh (Parkinson *et al.*, 1999b), the dopamine innervation of the accumbens (Taylor & Robbins, 1986; Cador *et al.*, 1991; Wolterink *et al.*, 1993), and the CeA (Robledo *et al.*, 1996), once again raising the possibility that the CeA normally plays a part in controlling Acb DA during appetitive Pavlovian tasks (see p. 43).

However, the glutamatergic inputs to the Acb involved in conditioned reinforcement appear to differ from those involved in autoshaping (for which projections from the ACC to the AcbC are critical). The efficacy of conditioned reinforcers is impaired by lesions of the BLA (Cador *et al.*, 1989; Everitt *et al.*, 1989; Burns *et al.*, 1993), but not the mPFC or the ventral subiculum (Burns *et al.*, 1993). Information of some sort about conditioned reinforcement must arrive at the Acb for its effects to be potentiated selectively by intra-Acb amphetamine — either the Acb must have *direct* information regarding the motivational significance of the CRf, or other structures that cause the animal to respond for CRf must provide the Acb with information about the identity of the current prepotent response, in order for intra-Acb amphetamine to potentiate this response selectively. Thus it appears that information regarding the conditioned value of the CS depends upon the BLA and is conveyed to the Acb (though not necessarily directly or exclusively), where its effects can be potentiated or ‘gain-amplified’ by dopamine (Robbins & Everitt, 1992). The BLA projects strongly to the Acb (both core and shell; Brog *et al.*, 1993), and while shell lesions abolish the effects of intra-Acb amphetamine, lesions of the core alter the normal response to intra-Acb amphetamine, such that amphetamine increases responding on both levers — a loss of response selectivity (Parkinson *et al.*, 1999c). Whether the orbitofrontal cortex or ACC also contributes to responding

for conditioned reinforcement is unknown at present; the role of the ACC is investigated directly in Chapter 3.

It remains a mystery as to precisely how the core and shell subdivisions of the Acb interact in this task. Apparently, information regarding the conditioned reinforcer arrives at the core and the shell (directly or indirectly from the BLA), but the ability of amphetamine to amplify the effects of this information depends upon the dopaminergic innervation of the Acb and the integrity of the shell, while the response selectivity of this amplification depends upon the core. Perhaps the enhancement of responding induced by intra-shell amphetamine is directed by the core towards the correct response. Though the core and shell do not project to each other directly (Brog *et al.*, 1993), the shell may modify the information passing through the core via indirect routes: Haber *et al.* (2000) have shown that the shell projects to regions of the VTA that innervate the shell itself, but also to VTA regions that project to the core; thus, the shell may exert control over dopamine function in the core. (Similarly, the core may be able to exert control over the dopamine projection to itself and to the central striatum, which may control the dorsolateral striatum in an ‘ascending spiral’ — a progression from limbic, through cognitive, to motor corticostriatal loops; see Figure 8.) Nevertheless, if this scheme is applicable, it is unclear why shell lesions block the effect of amphetamine injections into the core (Parkinson *et al.*, 1999b). Alternatively, it may be that intra-Acb amphetamine’s effects on the vigour and direction of behaviour (dependent upon the AcbSh and AcbC, respectively) are not integrated within the Acb, but are integrated at downstream sites (a possible candidate being the ventral pallidum; Fletcher *et al.*, 1998).




Figure 8. Organization of striatoni-grostriatal projections in the primate, illustrating one putative, dopaminergic mechanism by which corticostriatal loops influence each other in a hierarchy (Haber *et al.*, 2000). The colours illustrate the corticostriatal loops (*red*, limbic; *green*, associative; *blue*, motor). The AcbSh projects to regions of the VTA that innervate the AcbSh (*red*), but also the AcbC (*orange*). Similarly, the AcbC projects to areas of the VTA that innervate itself, but also to regions that project to the dorsomedial striatum. This spiral continues through more dorsal striatal regions (*yellow* → *green* → *blue*). The magnified oval region illustrates a hypothetical regulatory mechanism: striatal projections to those VTA neurons providing a closed-loop feedback projection terminate directly on the dopaminergic cell, inhibiting VTA neuron firing; however, striatal projections to those VTA neurons providing a feedforward projection to a different striatal region terminate on inhibitory interneurons, disinhibiting the dopaminergic innervation of the adjacent region (in this case, the AcbC). *DL-PFC*, dorsolateral prefrontal cortex; *IC*, internal capsule; *OMPFC*, orbital and medial prefrontal cortex; *S*, nucleus accumbens shell; *SNc*, substantia nigra pars compacta; *SNr*, substantia nigra pars reticulata; *VTA*, ventral tegmental area. Reproduced from Haber *et al.* (2000).

Surprisingly, animals with lesions of the AcbC or AcbSh retain the basic conditioned reinforcement effect — the ability to respond preferentially on a lever delivering an appetitive CS — even though their response to psychostimulants is altered. One possibility is that the expression of conditioned reinforcement itself does not depend entirely on Pavlovian processes. Clearly, Pavlovian conditioning is the mechanism by which a stimulus is established as a conditioned reinforcer. However, the *expression* of this learning might be through several mechanisms; in particular, it must be considered that the conditioned reinforcer becomes a true declarative instrumental ‘goal’ (discussed above, p. 31). Given that the accumbens is not necessary for animals to represent the value of an instrumental goal (Balleine & Killcross, 1994), it might not then be expected that Acb lesions would remove all effects of a conditioned reinforcer. The BLA is known to be important for the representation of the value of conditioned reinforcers (Cador *et al.*, 1989; Burns *et al.*, 1993; Killcross *et al.*, 1998); other candidate structures for the representation of the value governing instrumental responding for CRf are the insular cortex and orbitofrontal cortex (see pp. 44–46). In contrast to the results discussed above, however, Dix *et al.* (2000) recently reported that excitotoxic lesions of the whole Acb impaired the ability of rats to direct instrumental behaviour on the basis of conditioned punishment; it remains to be seen what the effects of selective AcbC and AcbSh lesions would be on this task. While this result might be interpreted as a difference between the circuits mediating the effects of appetitive and aversive CSs, it remains possible that both the AcbC and AcbSh contribute to the basic effect of conditioned reinforcement and that excitotoxic lesions of the whole Acb would impair responding for conditioned reinforcement, though dopamine depletion does not (Taylor & Robbins, 1986).

Pavlovian–instrumental transfer is impaired by lesions involving the nucleus accumbens

Conditioned reinforcement is a phenomenon by which a Pavlovian CS is delivered contingent upon responding. However, the accumbens is also critical for the impact of *noncontingent* Pavlovian conditioned stimuli. This has been demonstrated clearly by Pavlovian-to-instrumental transfer (PIT) experiments. If an animal is trained to press a lever for food and subsequently tested in extinction, presentation of a Pavlovian CS that predicts the same food increases the rate of lever-pressing (Estes, 1948; Lovibond, 1983). Lesions of the AcbC (Hall *et al.*, 1999) abolish PIT (see also de Borchgrave, 1995), as does systemic treatment with dopamine receptor antagonists (Smith & Dickinson, 1998; Dickinson *et al.*, 2000). A recent study also demonstrated that PIT can be enhanced by intra-accumbens amphetamine in the same way that conditioned reinforcement is. Wyvell & Berridge (2000) trained rats to respond on a lever for food, and also gave them associative pairings of a lever/light CS with that food. In a subsequent extinction test, they found that intra-Acb amphetamine (targeted at the AcbSh) increased the ability of the CS to potentiate responding, whether the CS was located within the lever (in which case the results might reflect a potentiation of autoshaping rather than PIT) or was a diffuse auditory stimulus. Finally, PIT is also impaired by CeA lesions (Hall *et al.*, 1999), leading Hall *et al.* (1999) to speculate that the ability of an appetitive Pavlovian CS to potentiate instrumental behaviour depends on the mesolimbic dopamine system, presumably under the control of the CeA (see pp. 43 & 47).

As described earlier, PIT can be subdivided into a general arousing effect of appetitive Pavlovian stimuli and a more informational component by which Pavlovian CSs selectively potentiate instrumental behaviour with which they share an outcome. It remains to be seen whether the arousing (general) and informational (specific) mechanisms by which noncontingent stimuli potentiate behaviour are the same as those involved for contingent stimuli (conditioned reinforcers). In both cases, such evidence as is available suggests that the informational component is subserved by glutamatergic projections from limbic structures such as the amygdala and ACC, with that information arriving directly or indirectly at the Acb, whilst the arousal component depends upon ascending projections from the isodendritic core to the Acb

(Taylor & Robbins, 1984; Cador *et al.*, 1989; Bussey *et al.*, 1997a; Han *et al.*, 1997; Hall *et al.*, 1999). Consider a response-specific PIT task. If PIT is truly comparable to the potentiation of responding for CRf by intra-Acb amphetamine, then one would expect the AcbC to be responsible for the response selectivity of PIT, and the AcbSh to be critical for the potentiation itself. If, on the other hand, response-specific PIT has a great deal in common with 'simple' PIT, one might expect AcbC lesions to abolish PIT entirely, as in the study of Hall *et al.* (1999). These predictions will be compared in Chapter 4 by testing the effects of lesions to the nucleus accumbens core and shell on response-specific PIT.

Contribution of the nucleus accumbens to complex behaviour

The role of these motivational processes in performance under different schedules of reinforcement is imprecisely understood. From an economic point of view, there is a high probability of executing an action when the motivation to perform that action exceeds the response costs, which include the work-related costs (effort). Schedule performance depends on these two variables; indeed, the progressive-ratio (PR) schedule is based on these principles. Salamone and colleagues have demonstrated that 6-OHDA-induced dopamine depletion of Acb impairs the ability of animals to overcome response costs (Salamone, 1994). Thus, DA-depleted rats will forgo the opportunity to press a lever for a preferred food, instead consuming more of a less-preferred but freely available food (Salamone *et al.*, 1991; Cousins *et al.*, 1993). Similarly, dopamine depletion impairs responding on high-rate but not on low-rate schedules (McCullough *et al.*, 1993; Salamone *et al.*, 1993; Sokolowski & Salamone, 1998; Aberman & Salamone, 1999). Of course, some of these results may be explained in terms of motoric impairments (such as a reduction in the maximum possible rate of responding). Cousins *et al.* (1996), however, tested rats in a T-maze in which one arm led to a large reward, but was obstructed by a barrier over which the rats had to climb, while the other arm, though it led to a small reward, was unobstructed. Cousins *et al.* found that while Acb DA depletion significantly reduced rats' preference for the arm that contained the barrier in this situation, DA depletion had minimal effects on rats' ability to climb the barrier when no alternative reward was available, suggesting that Acb DA depletion has effects that cannot be attributed purely to motor deficits. These results are compatible with the loss of a dopaminergic motivational influence that contributes to normal performance. Indeed, Acb dopamine depletion does not only impair responding under instrumental reinforcement schedules, but also displacement behaviour occurring when food is delivered on a fixed-time schedule (Robbins & Koob, 1980). Such behaviour cannot easily be described as carrying a response cost, whereas it may reflect a potentiation of irrelevant available behaviours by a motivational effect of the food (Robbins & Koob, 1980).

The interpretation that the Acb contributes Pavlovian conditioned motivation to behaviour is compatible with the view that it mediates aspects of preparatory behaviour, temporally distant from the goal of behaviour (as opposed to consummatory behaviour, temporally close to the goal). As an example of such a distinction, lever-pressing by male rats for access to a female has been doubly dissociated from unconditioned sexual behaviour (Everitt *et al.*, 1987; Everitt & Stacey, 1987). This distinction has been phrased in various ways — preparatory and consummatory (Blackburn *et al.*, 1987; Robbins & Everitt, 1992), seeking and taking (Arroyo *et al.*, 1998; Everitt *et al.*, 1999), and sign tracking and goal tracking (Hearst & Jenkins, 1974). Manipulations of the Acb, including 6-OHDA lesions and systemic injections of dopamine receptor antagonists, have been shown to reduce the preparatory aspects (including rate of responding) of behaviour directed towards both food and (in male rats) a sexually receptive female, whilst leaving consummatory behaviour unaffected (Blundell *et al.*, 1977; Koob *et al.*, 1978; Kelley & Stinus, 1985; Blackburn *et al.*, 1987; Everitt, 1990). Schedule-induced polydipsia (SIP), a phenomenon whereby excessive drinking is produced by the intermittent presentation of small amounts of food, is also disrupted

selectively by 6-OHDA lesions of the Acb, but not of the dorsal striatum (Robbins & Koob, 1980; Mitelman *et al.*, 1990). Acb lesions abolish SIP, leaving drinking/ingestion intact, whilst lesions of the dorsal striatum do not affect SIP but impair the ability of animals to drink effectively. In almost all paradigms studied, manipulations of limbic corticostriatal circuitry affect preparatory but not consummatory behaviour (Robbins & Everitt, 1992). The functional importance of Acb-dependent preparatory behaviour has also been demonstrated in a naturalistic setting by Whishaw & Kornelsen (1993). Rats normally carry food to a refuge to eat it, and when sated, carry the remaining food to hoard; rats with ibotenic acid or 6-OHDA lesions of the Acb were selectively impaired in this preparatory behaviour, failing to carry food to hoard it. The same rats were not impaired at carrying-to-eat, or eating itself.

Finally, a wide range of other tasks that depend on the effect of Pavlovian stimuli on instrumental or approach behaviour are also sensitive to lesions of the Acb or its afferents. The level of instrumental lever pressing is reduced by excitotoxic lesions of the Acb (Balleine & Killcross, 1994), consistent with the loss of a Pavlovian motivational effect that normally potentiates responding. Kelley *et al.* (1997) have also demonstrated a profound effect of intra-Acb infusions of glutamate receptor antagonists on Pavlovian and instrumental responding. Bilateral lesions of the BLA, or Acb, or a disconnection of the two, abolish a previously acquired conditioned place preference (CPP) for food (Everitt *et al.*, 1991); similarly, lesions of structures downstream from the Acb, including the ventral pallidum and mediodorsal thalamus, impair acquisition of a CPP (McAlonan *et al.*, 1993). The BLA and Acb are also critical for the acquisition of responding under second-order schedules of sexual or cocaine reinforcement (Everitt *et al.*, 1989; Whitelaw *et al.*, 1996), in which the second-order CS is critical for responding in normal animals (Arroyo *et al.*, 1998) (for similar studies using heroin reinforcement, see Robbins *et al.*, 2000; Alderson *et al.*, in press-a; Alderson *et al.*, in press-b). As discussed above (p. 42), lesions of the ACC, or AcbC, or a disconnection of the two, impair the acquisition of autoshaping (Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c).

Summary

The nucleus accumbens is a key site mediating the ability of a Pavlovian conditioned stimuli to invigorate and direct behaviour; it is critical for autoshaping, the effect of psychostimulant-amplified conditioned reinforcers on instrumental responding, and PIT. This motivational influence of Pavlovian CSs has been termed *incentive salience* (Robinson & Berridge, 1993; Berridge & Robinson, 1998), or 'Pavlovian incentive value' (Dickinson *et al.*, 2000), to distinguish it from the instrumental incentive value of Dickinson and colleagues (Dickinson, 1994; Dickinson & Balleine, 1994) (*q.v.* and see pp. 24–27 for discussion of the differences between the two).

DELAYED REINFORCEMENT

Delayed reinforcement is of interest from two theoretical perspectives. Firstly, how do animals succeed in bridging delays to reinforcement at all? Natural reinforcers always follow the action that obtained them by a delay, even if it is short. Thus, to control the world successfully, animals must be able to use delayed reinforcement. In some species, the delay to reinforcement may be very long indeed; humans routinely make decisions on the basis of outcomes that are decades away. Secondly, what accounts for differences in individuals' ability to choose delayed rewards? Why are some individuals impulsive in their choices? These questions will be considered in order.

Delayed reinforcement in learning

Early theorists considered the fundamental problem of delayed reinforcement: how a response can be strengthened by reinforcement that follows it. Hull (1932) postulated that the strength of an S–R association is inversely related to the delay between the response and the reinforcement, assuming a logarithmic relationship. Indeed, instrumental learning has repeatedly been shown to be a decreasing function of the delay (e.g. Lattal & Gleeson, 1990; Dickinson *et al.*, 1992). In several of the early studies, the delay was bridged by distinctive cues or environments. The cue that precedes eventual reward has the potential to become a secondary or conditioned reinforcer; thus the ‘underlying’ delay gradient function was unclear. In an effort to minimize the contribution of conditioned reinforcement, Grice (1948) trained rats on a visual discrimination task with delayed reinforcement. The rats had a choice of a white or a black start alley (which varied in their left/right position); the delay was provided by two grey alleys of variable length which terminated in two grey goal boxes. Choosing white led to a goal box with food; choosing black led to an empty box. Grice found that learning was noticeably impaired by as short a delay as 0.5 s, and severely impaired by 5 s. This deficit could be ameliorated by having more discriminable (black and white) goal boxes, or forcing the rats to make discriminable motor responses (climbing an incline or dodging between blocks) in the black and white start alleys.

Grice argued that Hull’s primary delay of reinforcement did not exist and that learning under conditions of delayed reward was due to immediate secondary reinforcement, based on traces of visual or proprioceptive stimuli. Clearly, if the primary gradient does exist, it is steep; the distinction becomes one of whether the delay applies to response reinforcement (Hull) or stimulus–reward association (Grice).

One other perspective deserves comment: that of Killeen & Fetterman (1988), who suggested that the very idea of a ‘delay gradient’ is misleading. In their model, reinforcement always strengthens the responses that the animal is presently making, and never acts ‘backwards in time’ to strengthen past responses. The observed ‘gradient’ stems from the fact that the animal has a finite probability of leaving the behavioural state it was in when it responded; if reinforcement follows immediately, there is a high probability of strengthening the response that caused reinforcement, but the longer the reinforcer is delayed, the greater the chance that the animal has moved to another state, in which case a different response will be reinforced. This point has also been made by Spence (1956), Mowrer (1960), and Revusky & Garcia (1970); see also Mackintosh (1974, pp. 155–159).

It is obviously impossible for response–reinforcement or stimulus–reinforcement learning to occur unless the trace of the response or the stimulus persists to be reinforced or associated. Whichever of the three perspectives has most merit, the point is made that small delays of reinforcement can markedly impair learning, that stimuli differentially associated with reward can improve this performance, and that interoceptive cues can sometimes perform this function.

Choice, and pathological choice, from the perspective of utility theory

Before considering the role of delayed reinforcement in choice behaviour, I will briefly review one theoretical approach to choice behaviour, utility theory, that has explicitly or implicitly underlain many studies using delayed reinforcement.

Utility theory

Formal utility theory is based on six axioms that define attributes of preference that perfectly rational agents should possess (von Neumann & Morgenstern, 1947) (reviewed by Russell & Norvig, 1995). One, for example, is *transitivity*: if an agent prefers *A* to *B* and *B* to *C*, then it must prefer *A* to *C*. If the agent

violated this principle, preferring $A > B > C > A$, then an observer could offer the agent C in exchange for A and a small monetary payment; similarly B for C and A for B , after which the agent ends up in its original state but with less money, which (assuming money is desirable) is irrational.

Given that a rational agent obeys these axioms, then there must exist a *utility function* U that assigns a real number to every outcome O such that $U(O_1) > U(O_2)$ if O_1 is preferred to O_2 , and $U(O_1) = U(O_2)$ if the agent is indifferent between the two outcomes.

Goal-directed action requires that the agent assigns value (goal status) to outcome states, but also that it knows the consequences of its actions. To allow for the fact that actions may not always have totally predictable consequences, the agent's knowledge about the causal nature of the world may be represented in the form $p(\text{action} \rightarrow \text{outcome}_n \mid \text{evidence})$ denoting the probability, given the available evidence, that *action* causes *outcome_n*. The *expected utility* of an action is therefore given by:

$$EU(\text{action} \mid \text{evidence}) = \sum_n p(\text{action} \rightarrow \text{outcome}_n \mid \text{evidence}) \cdot U(\text{outcome}_n)$$

Rational decision-making follows if the agent selects the action with the maximum expected utility (the MEU principle). The theory specifies neither the utility functions themselves — anything can be valued — nor the way that the decision is arrived at, which may be explicit or implicit.

The formal decision-making approach described suffers from two particular deficiencies. Firstly, computing the expected utilities takes finite time. It may often be better to make an imperfect decision quickly than *eventually* to make what *would* have been the perfect decision. In artificial intelligence, this has proved a difficult problem (Russell & Norvig, 1995). Secondly, the MEU principle implies that in identical situations, the same action will always be taken (it is a 'pure' strategy). However, game theory (von Neumann & Morgenstern, 1947) has shown that there are many situations involving choice under uncertainty when the optimal strategy is to assign probabilities to making different choices but to let the actual decision be governed by chance (a 'mixed' strategy). Even using this method, Gödel's (1931) incompleteness theorem implies that no concept of rationality can be optimal in every situation. How randomness is used in decision-making is poorly understood; Mérö (1998) provides an entertaining look at these issues.

Pathological choice in the context of utility theory

There are two ways within the framework of utility theory to produce 'pathological' decision-making. One is to alter the utility functions. For example, assigning a higher utility to poverty than wealth would cause a perfectly rational agent to give its money away; if gambling had intrinsic utility then an agent might gamble despite the financial loss. While the underlying choice remains rational, the agent's preferences generate abnormal behaviour. Indeed, some investigators see it as axiomatic that animals make rational or optimal decisions (see Williams, 1994b, pp. 91/94), so that the experimenter's job is to discover the value system of the subject.

The other mechanism, considered less often, is that utilities are computed normally but the decision-making process itself fails. Indeed, normal humans are not 'normative': they systematically deviate from the axioms of decision theory (Kahneman *et al.*, 1982; see also Chase *et al.*, 1998), which, incidentally, is why computerized systems can outperform human experts (Heckerman *et al.*, 1992).

The distinction is difficult. As an illustration, consider a smoker who desires to abstain but lights a cigarette. Are we to consider the decision flawed or the actual utility of smoking higher than he thought? If 'optimality can be considered axiomatic' (Williams, 1994b, p. 94), the latter is the case, but such a theory cannot distinguish between the act of our relapsing smoker and one who has no wish to give up. Nev-

ertheless, the distinction seems important; these questions only begin to make sense within a reductionist approach to the way the brain reaches decisions.

Views of choice in the brain

To choose between two goals that differ in nature, such as food *v.* money, they must be compared on a single dimension. Utility functions achieve this by converting multifactorial alternatives to real numbers. Neurally, a similar process is logically unavoidable — if at no earlier stage of processing, incompatible behaviours must compete for access to motor output structures (although there is no *a priori* reason why the neural comparison process should be simple or linear).

There is a long history of behavioural research into the computation of reward utility and consequent behavioural strategy (reviewed by Williams, 1994b), including the utility of artificial reinforcers (see Shizgal, 1997). One approach used was to calculate the efficacy of reinforcement by establishing the relationship between response *rate* and the frequency and amount of reinforcement; however, such attempts soon established that this relationship was not simple (see Williams, 1994b, pp. 82–83). For example, response rates are affected by whether a ratio or an interval schedule of reinforcement is used, even when the reinforcement rate is identical (Dawson & Dickinson, 1990). Similarly, the mechanisms governing motor aspects of responding are neurally dissociable from motivational mechanisms (see e.g. Robbins & Everitt, 1992).

The matching law and related research: ‘top-down’ analyses of choice

Another approach has been to relate reinforcement efficacy to choice behaviour. This literature stems from the discovery by Herrnstein (1961; 1970) of the ‘matching law’. Herrnstein (1961) trained pigeons to respond on two concurrent variable interval (VI) schedules, and varied the relative availability of reinforcement on the two schedules while holding the overall reinforcement rate constant. He observed that the proportion of the total behaviour allocated to each response key approximately matched the proportion of reinforcers allocated to that key. This defines the matching law:

$$\frac{R_1}{R_1 + R_2} = \frac{r_1}{r_1 + r_2}$$

where R represents the behavioural response rate for each alternative, and r the reinforcement. Herrnstein (1970) extended this relationship to take account of more than two alternatives, particularly including ‘unmeasured’ activities the animal may engage in, and derived a ‘general principle of response output’ (Herrnstein, 1970, p. 256):

$$R_1 = \frac{kr_1}{r_1 + r_e}$$

where R_1 is the rate of the response being measured, r_1 is the quantity of reinforcement for that response, r_e is the reinforcement for all other responses, and k is a parameter determining the maximum response rate. Although there are situations where the matching law is not useful — in particular, ratio schedules, where the distribution of reinforcement necessarily *follows* the distribution of responding — a vast literature has sought to define the effects of varying parameters of reinforcement (such as rate, probability, delay, and magnitude) based on this work (see de Villiers & Herrnstein, 1976).

Problems have emerged. In many circumstances, subjects have been found to ‘overmatch’ (exhibit preferences that are exaggerated relative to the predictions of the matching law) or ‘undermatch’ (exhibit reduced preferences), requiring further development of the mathematical models (Baum, 1974; Baum, 1979), though it has been argued that this is a circular approach (Rachlin, 1971). Maximum response rates (k in the equation above) have been shown to vary with the kind of reinforcement used (Belke, 1998),

violating an assumption of Herrnstein's law. Nevertheless, the matching law and its extensions do a good job of describing the relationship between reinforcement rate and behaviour on concurrent VI and concurrent-chain schedules (Williams, 1994b).

The matching law described a molar property of behaviour — that is, the overall distribution of a large number of responses. As responses are made on a moment-to-moment basis, the question arises of what 'molecular' choice process operates to produce matching at a molar level. Suggestions vary from 'momentary maximizing' theory, which suggests that subjects choose (in all-or-none fashion) the response with the highest instantaneous reinforcement probability, to the idea that matching is the basic choice rule (see Mackintosh, 1974, pp. 192–195; Williams, 1994b).

Relating choice to 'value'

All these theories share a theoretical basis: it is assumed that some value is computed for each alternative behaviour, and a single decision rule allocates behaviour according to the relative distribution of values. In order to produce a single value for each alternative, different reinforcement parameters (rate, magnitude, delay, and probability) converge on a single dimension (Baum & Rachlin, 1969). Often, the effects of these different parameters are assumed to be calculated independently (Killeen, 1972; Rachlin *et al.*, 1991; Ho *et al.*, 1999). Though some investigators have supported the latter assumption (Mazur, 1987; Mazur, 1997), using different techniques, others have found that the effects of reinforcer delay and magnitude are not independent (Ito, 1985; White & Pipe, 1987). In either case, the assumption that all choice alternatives are reduced to a single value and then compared in order to select the option with the greatest value corresponds directly to a form of utility theory, as described above.

Fragmenting choice: a neuropsychological or 'bottom-up' approach

We have seen how utility theory can fail to characterize human decision-making (Kahneman *et al.*, 1982), just as similar approaches have not fully characterized choice in other animals (Williams, 1994b, p. 105). Perhaps more success can be achieved by considering the multiple psychological systems that have been discovered to contribute to instrumental performance. In this framework, behaviour and choice are seen as the asymptotic sum of contributions from cognitive goal-directed systems, habitual responding and other motivational influences (e.g. Dickinson, 1994). As we have seen, rats possess *at least* two representations of the value of foodstuffs (Dickinson & Balleine, 1994), namely hedonic value and the incentive value governing instrumental responding; Pavlovian incentive value is probably a third (see pp. 24 & 28). An analysis of the neuropsychological mechanisms by which these multiple motivational systems calculate the value of environmental events and interact with each other may prove more productive than the 'top-down' approach. To take a hypothetical example, suppose that stimulus–response habits obey the matching law, but that cognitive, voluntary decisions can override habits in some circumstances and have a different value system. It is likely that acknowledging the existence of these two systems, and determining when each operates, will more rapidly lead to an accurate description of choice behaviour than attempting to model choice with a single, but highly complicated, value system.

Neuropsychological research along these lines is a young field. As has been outlined above, consideration of the neural basis of Pavlovian and instrumental conditioning in animals has led to the identification of several brain regions and neurotransmitter systems that are involved in reinforcement and value assessment (including the reinforcing effects of drugs of abuse). This literature supports neuropsychological data derived from studies of humans with acquired disorders of decision-making. Noteworthy among these are studies of humans with damage to the orbitofrontal cortex or amygdala (including the famous case of orbitofrontal cortex damage in Phineas Gage, first reported by Harlow, 1868), who exhibit im-

paired choice behaviour despite apparent knowledge that they are choosing poorly (Damasio, 1994; Bechara *et al.*, 1998). Following the pioneering theories of Nauta (1971), this work has led to the development of a specific theory of a process contributing to choice, namely the somatic marker hypothesis (Damasio, 1994; Damasio, 1996; Bechara *et al.*, 2000). This theory proposes the existence of a non-conscious, rapidly-retrieved utility signal that improves decision-making performance by removing poor options from the consideration of a computationally-intensive cognitive process. These signals appear to have a measurable autonomic correlate in galvanic skin responses (Bechara *et al.*, 1996; Bechara *et al.*, 1997) and depend upon the integrity of the amygdala and orbitofrontal cortex (Bechara *et al.*, 1998; Bechara *et al.*, 1999). This may represent the best-characterized neural correlate of decision-making.

The other major avenue of investigation into pathological decision-making has concentrated on the phenomenon of impulsivity. Research into impulsivity is deeply interwoven with the study of delayed reinforcement; these areas will be reviewed next.

Impulsivity and impulsive choice

Impulsivity was well known to the ancient Greeks. The character flaw *akrasia* (weakness of will, lack of self-control, or incontinence) is a deficiency of the power to act as one judges best in the face of competing motivation. Aristotle saw it as a commonplace deviation from the norm of men:

‘The incontinent man, knowing that what he does is bad, does it as a result of passion, while the continent man, knowing that his appetites are bad, refuses on account of his rational principle to follow them.’

Nicomachean Ethics (Aristotle, 350 BC / 1925, book 7, chapter 1)

‘It is plain, then, that incontinent people must be said to be in a similar condition to men asleep, mad, or drunk.’

(Book 7, chapter 3.)

‘Now incontinence and continence are concerned with that which is in excess of the state characteristic of most men; for the continent man abides by his resolutions more and the incontinent man less than most men can.’

(Book 7, chapter 10.)

Aristotle’s definition of incontinence focuses on an inability to suppress one’s desires in favour of more rational, high-minded resolutions. However, the term ‘impulsivity’ has been applied to many different aspects of maladaptive choice. Ainslie (1975) summarized three guesses about why humans are prone to obey ‘impulses’:

- (1) that they lack insight into the consequences of their actions — a defect in instrumental contingency learning;
- (2) they are aware of the consequences of their actions, but are unable to suppress ‘some lower principle (the devil, repetition compulsion, classical conditioning)’ — a defect in response inhibition;
- (3) they are aware of the consequences of their actions, and choose rationally according to their value system, but their values are ‘innately distorted so that imminent consequences have a greater weight than remote ones’ — reduced value of delayed reinforcement.

Impulsivity may be given a broader scope still; Evenden's (1999b) review of the field encompasses all the above and adds 'preparation impulsivity' (reaching a decision before adequate information is gathered) and 'execution impulsivity' (interrupting a chain of behaviour before its goal is achieved) (Evenden, 1998, p. 37). Critically, these aspects of impulsivity may be dissociated pharmacologically, implying that they reflect genuinely different underlying processes (Evenden, 1999b).

Impulsivity may be considered a normal personality trait (Eysenck & Eysenck, 1977; Barratt & Patton, 1983; Eysenck, 1993), as Aristotle did, but it is also a feature of a number of clinical disorders. These include personality disorders (antisocial personality disorder and borderline personality disorder; APA, 1994), impulse control disorders, including drug addiction ('substance abuse disorder'; APA, 1994), and attention-deficit/hyperactivity disorder (ADHD), a prevalent disease of childhood of which impulsivity is one sign (Evenden, 1998; Sagvolden & Sergeant, 1998).

Multifaceted though impulsivity is, the present thesis will focus exclusively on *impulsive choice*, exemplified by the inability of an individual to choose a large delayed reward in preference to a small immediate reward (or, in an aversive context, an inability to choose a small immediate penalty in preference to a large delayed penalty). This form of impulsivity can be characterized as pathological hypersensitivity to delays of reinforcement (though the formally identical aversive analogy may make it clearer that impulsive choice may also have something to do with relative insensitivity to differences in reinforcer magnitude). What, though, is 'normal' sensitivity to delayed reward?

Delayed reinforcement in choice

In a typical situation, a subject chooses between an immediate, small reward or a large, delayed reward; the time discounting function quantifies the effect of the delay on preference. Kacelnik (1997) points out that economic models of choice tend to be based on exponential time discounting functions. If the starting assumption is that delayed reward is preferred less because there is a constant probability of losing the reward per unit of waiting time, or that there is a constant 'interest rate' for the reward obtained immediately (and that the subject's behaviour is attuned to this fact, i.e. that choice is normative) then exponential models emerge. If a delayed reward of magnitude A is chosen and there is a probability p of loss in every unit of time waited, the perceived value V of the delayed reward should be $V = A(1 - p)^T = Ae^{-kT}$ where $k = -\ln(1 - p)$.

However, the exponential model has been emphatically rejected by experimental work with humans and other animals. The literature on human cognitive decisions will not be considered here. The rat literature contains several demonstrations (many based on the adjusting-delay task of Mazur, 1987) procedure, using natural reinforcers and intracranial self-stimulation (or 'brain-stimulation reward') (Grice, 1948; Mazur, 1987; Mazur *et al.*, 1987; Richards *et al.*, 1997b), that time discounting is described well by a *hyperbolic* discount function (Figure 9) or at least a very similar power law (Grace, 1996). Kacelnik (1997) offers some explanations as to why hyperbolic discounting may be in some sense optimal. One interesting prediction from this function is that preference between a large and a small reward should be observed to reverse depending on the time that the choice is made (Figure 10), and such preference reversal is a reliable experimental finding (for references see Bradshaw & Szabadi, 1992). Of course, the neuropsychological system responsible for hyperbolic discounting is unknown — such discounting might, for example, result from poor knowledge of the action–outcome contingency at long delays, from weak stimulus–response habits, or from reduced utility of delayed rewards in the context of perfect contingency knowledge.




Figure 9. Hyperbolic discounting, governed by the equation

$$V = \frac{\text{magnitude}}{1 + K \cdot \text{delay}}$$

Large values of K give the steepest curve.




Figure 10. Preference reversal. If given a choice between an early reward of value 0.6 and a later reward of value 1, hyperbolic discounting predicts that the larger reward will be chosen if the choice is made far in advance (towards the left of the graph). However, as time advances, there comes a time just before delivery of the small reward when preference reverses and the small reward is chosen. Figure adapted from Ainslie (1975).

Neurochemical and neuroanatomical studies of delayed reinforcement

Serotonin (5-HT)

Abnormalities of the utility function for delayed reinforcement have been suggested to occur following neurochemical manipulations. The suggestion that serotonin is involved in impulse control follows from the twin observations that drugs that suppress 5-HT function appear to reduce behavioural inhibition, making animals more impulsive in the ‘motor’ sense (Soubrié, 1986), and that low levels of serotonin metabolites in cerebrospinal fluid are associated with impulsive aggression and violence in humans (e.g. Åsberg *et al.*, 1976; Linnoila *et al.*, 1983; Brown & Linnoila, 1990; Linnoila *et al.*, 1993) and risk-taking behaviour in monkeys (Mehlman *et al.*, 1994; see also Evenden, 1998). In the sphere of delayed reinforcement, forebrain serotonin depletion, which leads to ‘impulsive choice’ in a variety of paradigms (Wogar *et al.*, 1993b; Richards & Seiden, 1995; Bizot *et al.*, 1999), has been suggested to reflect a modification of the temporal discounting function (Wogar *et al.*, 1993b; Ho *et al.*, 1999). Specifically, 5-HT depletion is suggested to steepen the function, such that delayed rewards lose their capacity to motivate or reinforce behaviour. The animal becomes hypersensitive to delays (or hyposensitive to delayed reward). As delayed rewards have unusually low utility, the animal consistently chooses small, immediate rewards over large, delayed rewards, a characteristic of impulsivity (Ainslie, 1975). The specific contribution of different 5-HT receptor subtypes to choice of delayed reward has also been studied (Evenden & Ryan,

1996; Evenden, 1998; Bizot *et al.*, 1999; Evenden, 1999b; Evenden & Ryan, 1999), but this topic will not be pursued in detail.

Dopamine and attention-deficit/hyperactivity disorder (ADHD)

Much of the interest in the relationship between dopamine and impulsivity comes from the discovery that amphetamine and similar psychostimulants are an effective therapy for ADHD (Bradley, 1937). Though these drugs have many actions, they are powerful releasers of dopamine from storage vesicles in the terminals of dopaminergic neurons, and prevent dopamine re-uptake from the synaptic cleft, potentiating its action (for references see Feldman *et al.*, 1997, pp. 293/552/558). Sagvolden & Sergeant have proposed that many features of ADHD, including preference for immediate reinforcement and hyperactivity on simple reinforcement schedules (due to short inter-response times; Sagvolden *et al.*, 1998), are due to an abnormally short and steep delay gradient and that this is due to a hypofunctional dopamine system. Indeed, they go on to suggest Acb DA as the specific culprit (Sagvolden *et al.*, 1998; Sagvolden & Sergeant, 1998). Clearly, accumbens dopamine is implicated in aspects of responding for reinforcement, as discussed earlier, though its role is not yet fully understood.

Many of the inferences regarding the neural abnormalities in children with ADHD have in fact been drawn from studies of the spontaneously hypertensive rat (SHR), an inbred strain of rat that serves as an animal model of ADHD (Wultz *et al.*, 1990; Sagvolden *et al.*, 1992; Sagvolden *et al.*, 1993; Sagvolden, 2000). This rat exhibits pervasive hyperactivity and attention problems that resemble ADHD, is abnormally sensitive to immediate reinforcement in the sense that it exhibits a steeper ‘scallop’ of responding on fixed-interval (FI) schedules (Sagvolden *et al.*, 1992), and is impulsive on measures of ‘execution impulsivity’ (Evenden & Meyerson, 1999).

Examination of the brains of SHRs supports the assertion that they have an abnormality of dopamine systems. Depolarization- and psychostimulant-induced dopamine release in nucleus accumbens brain slices is altered in the SHR compared to Wistar Kyoto (WKY) progenitor control rats in a complex pattern that has been attributed to hypofunction of the mesolimbic dopamine system (de Villiers *et al.*, 1995; Russell *et al.*, 1998; Russell, 2000), though abnormalities have also been found in dopamine release in slices of dorsal striatum and prefrontal cortex (Russell *et al.*, 1995). Amygdala dysfunction has also been suggested (Papa *et al.*, 2000). Within the Acb, differences in gene expression and dopamine receptor density have been observed in both the core and shell subregions (Papa *et al.*, 1996; Carey *et al.*, 1998; Papa *et al.*, 1998).

Systemic psychopharmacological studies using normal animals provide an additional source of evidence regarding dopamine systems and choice of delayed reward. Many of these studies have examined the role of psychostimulant drugs such as amphetamine and methylphenidate, given these drugs’ efficacy in treating some symptoms of ADHD. However, conflicting results have been obtained in animal models. For example, psychostimulants have sometimes been found to promote choice of delayed rewards, and sometimes to impair it, in normal rats (Sagvolden *et al.*, 1992; Charrier & Thiébot, 1996; Evenden & Ryan, 1996; Richards *et al.*, 1997a; Richards *et al.*, 1999; Wade *et al.*, 2000). These apparent inconsistencies will be addressed in Chapter 6 by considering the potential contribution of conditioned reinforcement to the effects of psychostimulants in altering preference for delayed reward.

The prospect of delineating neural circuitry involved in choice of delayed reward

Although there has been considerable research on the neurochemical basis of tolerance to delayed rewards, together with correlative studies of cortical functional abnormalities in ADHD children and of regional differences in neurotransmission in the SHR, there have been few direct investigations of the role

of anatomically-defined brain structures in the capacity of animals to choose a delayed reward. Dopamine and serotonin affect this capacity, but where do they have their action? Abnormal functioning of prefrontal cortical regions, including the ACC, has been observed in humans with ADHD (Ernst *et al.*, 1998; Bush *et al.*, 1999), but it is not known whether these abnormalities are related to impulsive choice. In Chapter 7, I will consider the role of the ACC, the mPFC, and the AcbC, three structures that play a role in reinforcement processes, as outlined above, and that also receive serotonergic and dopaminergic projections (see Fallon & Loughlin, 1995; Halliday *et al.*, 1995).

Tasks used to study choice of delayed reinforcement

Two main approaches have been used to study delayed reinforcement: free-operant and discrete-trial procedures.

Free-operant tasks are typified by the concurrent-chain schedule (Autor, 1969; see Davison, 1987). In the most common variant, subjects respond on two concurrent VI schedules (the ‘initial links’). When one of the two schedules is completed, the other manipulandum is deactivated and the chosen schedule enters the ‘terminal link’ (in which reinforcement is provided on another schedule — for example, a fixed time schedule, in which noncontingent reinforcement is given after a fixed delay). Relative response allocation in the initial links is taken as a measure of relative preference for the two terminal link schedules. Versions of concurrent chain schedules that depend on subjects’ timing behaviour have also been developed (Gibbon & Church, 1981).

Though they allow accurate determination of relative response rates and, by inference, relative preference, such free-operant schedules carry two problems of interpretation. One is that the delays between initial-link responses and initiation of the terminal link may also form part of the delay to reinforcement; such delays are difficult to control for, and their importance may vary with the relative durations of initial and terminal links. The other applies to pharmacological and lesion studies: manipulations that affect an animal’s ability to produce motor responses, to switch between two responses, and to time their motor output, may all confound interpretation of the results (Ho *et al.*, 1999). For these reasons, free-operant schedules were not used to assess preference for delayed reward in the present thesis.

Discrete-trial schedules may also be divided into two classes. In the simplest type, the subject chooses between two mutually exclusive alternatives and preference is measured as the proportion of trials on which each alternative is chosen (e.g. Bradshaw & Szabadi, 1992; Evenden & Ryan, 1996). Choice in a T-maze has been used similarly (e.g. Bizot *et al.*, 1999). Although such schedules may not provide as accurate a measure of preference, as subjects tend towards exclusive preference on discrete-trial and ratio schedules (see Mackintosh, 1974, pp. 190–195), the response–reinforcer delay can be accurately controlled and instantaneous choice is free of the confounds discussed above regarding response rates and switching rates.

An alternative type of discrete-trial task, the adjusting-delay schedule, was invented by Mazur (1987). In this task, subjects choose in discrete trials between a *fixed* alternative, such as a small immediate reward, and an *adjusting* alternative, such as a larger reward delivered after a delay. This delay can alter. If the subject prefers the larger reward, the delay is lengthened, and if it prefers the smaller alternative, the delay is shortened, in an attempt to titrate the subject’s preference towards indifference. At this indifference point, the length of the adjusting delay is taken as a measure of the subject’s preference between the two reinforcers — it is the delay that ‘balances’ the difference in the magnitudes of the two reinforcers. (A similar schedule in which the amount of reinforcement is adjusted was recently described by Richards *et al.*, 1997b). The key advantage of the indifference-point methodology is that it allows quantitative es-

timation of subjects' preferences without assuming a particular relationship between reinforcer value and behavioural output — the only assumption required is that when two reinforcers are of equal value, behaviour is equally distributed between the two response alternatives. The technique also allows subjects' sensitivity to reinforcer delay to be distinguished from sensitivity to reinforcer magnitude (see Ho *et al.*, 1999). However, the schedule design is complex: not only does the delay affect the subject's choice, but choice affects the adjusting delay. Versions of this schedule have been used to assess the effects of chronic neurochemical manipulations (Wogar *et al.*, 1992; Wogar *et al.*, 1993b; Ho *et al.*, 1997). In Chapter 5, a group of rats are tested on a version of the adjusting-delay schedule to see if the task is suitable for other neurotoxic lesion and acute pharmacological studies.

It is worth noting at this point that there is a great difference between measuring delay preference using tasks in which trials occur at a fixed frequency (or at least, in which the subject's choice does not influence the time to the next choice-point) and those where choice can influence this frequency. In the latter kind of task, the subject may choose the small immediate reinforcer and have the opportunity to do so again very rapidly, so may be able to accumulate more reward by repetitive choice of the small reinforcer than by choosing the large one. If the overall trial frequency is held constant, however, the strategy that maximizes reward is always to choose the larger reinforcer; in this case, failure to do so can be attributed to its delay. Some studies of impulsive choice in children with ADHD used the former type of delay-of-gratification task, finding impairments, while comparable studies with a fixed trial length have failed to find differences (see Sonuga-Barke *et al.*, 1998). All studies reported in this thesis used a fixed trial frequency.

Conditioned reinforcement in choice of delayed rewards

Finally, when considering preference for delayed reinforcement, the role of conditioned reinforcers must be considered. Not infrequently, delay-of-reinforcement procedures have been used as a tool to study conditioned reinforcement; a stimulus is presented during the delay to reinforcement in the expectation that it will become a conditioned reinforcer. Such stimuli certainly affect choice behaviour (see Lattal, 1987), and tasks of this sort have served as the basis for attempts to quantify the value of conditioned reinforcers (Autor, 1969; Mazur, 1991; Mazur, 1995; Mazur, 1997), though the issues are complex (Williams, 1994a).

For the present thesis, it suffices to note that conditioned reinforcement can be an important factor influencing preference for delayed reward. In Chapters 6 and 7, I will seek to clarify the effects of certain drugs, including amphetamine, that are known to affect responding for conditioned reinforcement. In these experiments, explicit comparison will be made between the situation in which a stimulus is present during the delay to reinforcement and the situation in which no such stimulus is present. In Chapter 7, when lesion studies are conducted, the tasks used will not present a stimulus during this delay, to avoid this potential problem of interpretation.

ORGANIZATION OF EXPERIMENTAL WORK IN THIS THESIS

The experiments described in this thesis may be divided into two parts.

In Part 1 (Chapters 3 & 4), a clearer understanding is sought of the role of the ACC and Acb in basic Pavlovian and instrumental processes. In Chapter 3, I will consider in detail the functions of the ACC, as elucidated by previous rodent and primate studies, before investigating its contribution to simple Pavlovian conditioned approach, conditioned reinforcement, Pavlovian-instrumental transfer, and other simple

conditioning procedures. In Chapter 4, the involvement of the core and shell subdivisions of the Acb in Pavlovian–instrumental transfer is investigated, using a task with greater behavioural specificity than those previously used for this purpose.

In Part 2 (Chapters 5–7), an attempt is made to understand the contributions of limbic corticostriatal circuitry to the capacity of rats to choose delayed reward. In Chapter 5, a detailed examination is made of rats' performance on the adjusting-delay schedule of Mazur (1987), and this schedule is found unsuitable for pharmacological or lesion studies. In Chapter 6, a different task is developed, from that of Evenden & Ryan (1996). This task is used to investigate the effects of systemic dopaminergic drugs on choice of delayed reinforcement, and their interactions with stimuli that 'bridge' the delay to reinforcement. In Chapter 7, the same task is used to investigate the effects of destroying key elements of the limbic corticostriatal circuit — the ACC, mPFC, and AcbC — on choice of delayed reward, and the effects of intra-accumbens injections of amphetamine.

Chapter 2.

General methods

SUBJECTS AND HOUSING CONDITIONS

Subjects were male Lister hooded rats (Harlan-Olac UK Ltd) housed in a temperature-controlled room (minimum 22°C) under a 12:12 h reversed light–dark cycle (before March 2000, lights off 08:30 to 20:30; after March 2000, lights off 07:30 to 19:30). Subjects were approximately 15 weeks old on arrival at the laboratory and were given a minimum of a week to acclimatize, with free access to food, before experiments began. Experiments took place between 09:00 and 23:00, with individual subjects being tested at a consistent time of day. Unless otherwise stated, subjects were experimentally naïve, housed in pairs, provided with free access to water, and maintained throughout the experiment at 85–90% of their free-feeding mass using a restricted feeding regimen. Feeding occurred in the home cages at the end of the experimental day. In behavioural tasks where a significant amount of food was provided in the experimental chambers, the control software reported the amount of food delivered and this was used to correct the amount of food given in the home cages. All experimental procedures were subject to UK Home Office approval (Project Licences PPL 80/00684 and PPL 80/1324).

SURGERY

General surgical technique

Animals were anaesthetized with Avertin (2% w/v 2,2,2-tribromoethanol, 1% w/v 2-methylbutan-2-ol, also known as tertiary amyl alcohol, and 8% v/v ethanol in phosphate-buffered saline, sterilized by filtration, 10 ml/kg intraperitoneally)² and placed in a Kopf or Stoelting stereotaxic frame (David Kopf Instruments, Tujunga, California, USA; Stoelting Co., Wood Dale, Illinois, USA) fitted with atraumatic ear bars (Figure 11). The skull was exposed and a dental drill was used to remove the bone directly above the injection and cannulation sites. The dura mater was broken with the tip of a hypodermic needle, avoiding damage to underlying venous sinuses such as the superior sagittal sinus. Lesions and cannulation were accomplished according to the atlas of Paxinos & Watson (1996) using bregma as the origin and with the incisor bar set at 3.3 mm below the interaural line.

At the end of the operation, animals were given 15 ml/kg of sterile 5% w/v glucose, 0.9% w/v sodium chloride intraperitoneally. They were then left for 7 days to recover, with free access to food, and were handled regularly. Any instances of post-operative constipation were treated with liquid paraffin orally and rectally. At the end of this period, food restriction was resumed.

Excitotoxic lesions

Fibre-sparing excitotoxic lesions were made with quinolinic acid (Sigma, UK) dissolved in 0.1 M phosphate buffer (PB; composition 0.07 M Na₂HPO₄, 0.028 M NaH₂PO₄ in double-distilled water, sterilized by filtration) to a concentration of 0.09 M and adjusted with NaOH to a final pH of 7.2–7.4, or with 0.06 M ibotenic acid (Sigma, UK) in the same vehicle. Toxin was infused using one of two systems:

- a) through a 28-gauge stainless steel cannula (Ø 0.36 mm external, 0.18 mm internal; model C313, Plastics One, Roanoke, Illinois, USA, via Semat Technical Ltd, St Albans, UK) attached via polyethylene tubing to a

² Concentrations given as percentages are calculated as follows. A 1% solution, volume per unit volume (v/v), is a solution in which $1/100$ of the total volume is solute. A 1% solution, weight by unit weight (w/w), is one in which 1% of the total weight of the solution is solute; thus, a 1% solution implies 1 g of solute dissolved in 99 g of solvent. A 1% solution, weight by unit volume (w/v), is a solution of 1 g in a total volume of 100 ml (10 g l⁻¹); '100%' denotes 1 kg l⁻¹. Similarly, the notation '1:1000' denotes 1 g l⁻¹ (1 mg ml⁻¹).




Figure 11. Stereotaxic frame, from Carlson (1991).

10- μ l syringe (Hamilton Bonaduz AG, Bonaduz, Switzerland) mounted on a Harvard Apparatus (Edenbridge, UK) infusion pump;

- b) through a micropipette, manufactured using a Stoelting (Illinois, USA) pipette puller (model 52500) from glass pipettes (\varnothing 1.2 mm external, 0.69 mm internal; Intracel Ltd, Royston, UK) to give a final tip diameter of 50–100 μ m. The micropipette was glued to the tip of a 1- μ l syringe (SGE Ltd, Milton Keynes, UK) using Araldite epoxy resin (Ciba, UK), and the syringe was hand-driven during infusions.

Infusion times are shown in Table 7. At each site the injector was left in place for a specified time following infusion (see table) to allow diffusion away from the injection site and so to minimize ‘tracking’ of the toxin along the path of the cannula. Sham lesions were made in the same manner except that vehicle was infused.

Table 7. Lesion coordinates. Dorsoventral coordinates are either from dura (D) or skull surface at bregma (SS). Along the anteroposterior (AP), mediolateral (ML), and dorsoventral (DV) axes, positive coordinates are in the anterior, left, and superior directions respectively.

Lesion	Toxin and delivery system	Sites per side	AP	ML	DV	Volume per site	Infusion time	Diffusion time
Anterior cingulate cortex (peri-genual lesion)	quinolinic acid, 0.09 M, via cannula	6	+1.2	± 0.5	-3.0	0.5 μ l	1 min	1 min (lower sites); 2 min (upper sites)
			+1.2	± 0.5	-2.2			
			+0.5	± 0.5	-2.8			
			+0.5	± 0.5	-2.0			
			-0.2	± 0.5	-2.5			
			-0.2	± 0.5	-2.0 (SS)			
Nucleus accumbens core	quinolinic acid, 0.09 M, via micropipette	1	+1.2	± 1.8	-7.1 (SS)	0.5 μ l	3 min	2 min
Nucleus accumbens shell	ibotenic acid, 0.06 M, via micropipette	3	+1.6	± 1.1	-6.4	0.1 μ l	1 min	2 min
			+1.6	± 1.1	-6.9	0.1 μ l	1 min	1 min
			+1.6	± 1.1	-7.9	0.2 μ l	2 min	1 min
					(SS)			
Medial prefrontal cortex	quinolinic acid, 0.09 M, via cannula	4	+3.8	± 0.5	-1.5	0.5 μ l	1 min	2 min
			+3.3	± 0.5	-3.0			
			+3.3	± 0.5	-1.5			
			+2.6	± 0.5	-1.5			
					(D)			

Implantation of bilateral intracranial cannulae

Holes were drilled in the skull as described above. Four stainless steel screws were placed on each side about the burr holes and a pair of 22-gauge, bevelled stainless steel guide cannulae (13.0 mm long, \varnothing 0.7 mm external, \sim 0.4 mm internal; Coopers Needle Works, Birmingham, UK) were simultaneously lowered to the target position. Coordinates are given in Table 8, and a schematic is shown in Figure 12. With the cannulae held in place by inserters made of wire (\varnothing 0.36 mm) affixed to a steel frame, dental cement was applied around the cannulae and screws and allowed to dry. The inserters were then removed and the guide cannulae closed with stainless steel wire occluders (\varnothing 0.36 mm).

Table 8. Cannulae coordinates and intracerebral infusion parameters. First dorsoventral coordinate represents guide cannula tip location. Second coordinate represents injector tip location. Dorsoventral coordinates are from dura (D).

Cannulation	Drug	Sites per side	AP	ML	DV	Volume per site	Infusion time	Diffusion time
Nucleus accumbens	amphetamine sulphate, 0–20 μ g	1	+1.6	\pm 1.5	–5.0 –7.0 (D)	1 μ l	1 min	2 min




Figure 12. Intended location of guide cannulae (light grey) and injectors (dark grey), with the injector tip within the nucleus accumbens (1.6 mm anterior to bregma). The cannulae and injectors are drawn to scale, although the plastic injector mounting has been schematized. Brain section taken from Paxinos & Watson (1998).

HISTOLOGICAL ASSESSMENT

Perfusion and tissue fixation

At the end of experiments involving excitotoxic lesions or intracranial infusions, animals were deeply anaesthetized with Euthatal (pentobarbitone sodium, 200 mg/ml, minimum of 1.5 ml i.p.) and perfused transcardially with 0.01 M phosphate-buffered saline (PBS; composition 6.4 mM Na₂HPO₄, 1.47 mM NaH₂PO₄, 0.13 M NaCl, 2.68 mM KCl; the pK_a for this phosphate acid–base pair is 6.865, giving an approximate pH of 7.5) followed by 4% paraformaldehyde in PBS. Their brains were removed and postfixed in paraformaldehyde before being dehydrated in 20% sucrose for cryoprotection.

Nissl staining with cresyl violet

Cresyl violet stains for Nissl substance, the basophilic rough endoplasmic reticulum present in cytoplasm.

The brains were sectioned coronally at 60 µm thickness on a freezing microtome and every third section mounted on chrome alum (chromium potassium sulphate)/gelatin-coated glass microscope slides and allowed to dry. Sections were passed through a series of ethanol solutions of descending concentration (3 minutes in each of 100%, 95%, and 70% v/v ethanol in water) and stained for ~5 min with cresyl violet. The stain comprises 0.05% w/v aqueous cresyl violet (Raymond A. Lamb Ltd, Eastbourne, UK), 2 mM acetic acid, and 5 mM formic acid in water. Following staining, sections were rinsed in water and 70% ethanol before being differentiated in 95% ethanol. Finally, they were dehydrated and delipidated in 100% ethanol and HistoClear (National Diagnostics, UK) before being cover-slipped using DePeX mounting medium (BDH, UK) and allowed to dry.

The sections were used to verify cannula and lesion placement and assess the extent of lesion-induced neuronal loss. Lesions were detectable as the absence of visible neurons (cell bodies of the order of 100 µm in diameter with a characteristic shape), often associated with a degree of tissue collapse (sometimes with consequent ventricular expansion when the lesion was adjacent to a ventricle) and gliosis (visible as the presence of smaller, densely-staining cells).

Immunocytochemical staining for neuronal nuclei

Direct visualization of the location of neuronal nuclei was achieved using the NeuN antibody (Mullen *et al.*, 1992).

The immunocytochemical procedure may be summarized as follows. A primary immunoglobulin-G (IgG) antibody is raised in a mammalian species against a specific target of interest. The tissue being investigated is incubated with the primary antibody, and then unbound primary antibody is washed off. Next, the sections are incubated in secondary antibody. The secondary antibody was raised in a second mammalian species against the constant (Fc) portion of the IgG molecule of the first species; it therefore binds to all IgG antibodies of the first species, including the primary antibody. The secondary antibody has biotin attached to it. To this is added a complex of avidin and biotinylated horseradish peroxidase (HRP). Avidin is a protein with multiple biotin binding sites; thus, it can 'bridge' the biotinylated secondary antibody to the biotinylated HRP, ultimately binding HRP to the primary antibody and to the target labelled by it. The location of HRP can be visualized because HRP is an enzyme that catalyses the oxidation by hydrogen peroxide of various substrates, including the chromagen diaminobenzidine (DAB); the DAB oxidation product is a visible, insoluble precipitate.

Following cutting at 40 µm, sections were washed in 0.01 M phosphate-buffered saline (PBS; pH 7.4) and placed into primary antibody (NeuN monoclonal mouse anti-neuronal nuclear protein; Chemicon International Ltd, Harrow, UK), and incubated overnight with gentle agitation. The primary antibody was made up at a concentration of 1:10,000 in 0.01 M PBS containing the non-ionic detergent 0.4% t-octylphenoxypolyethoxyethanol (Triton[®] X-100) to solubilize the protein. The next day, the sections were washed three times in 0.01 M PBS over 30 minutes and incubated in secondary antibody for 90 min on a rotary shaker; the secondary antibody was biotinylated rabbit anti-mouse IgG (Vectastain ABC Kit, Vector, Burlingame, CA) at 1:200 in 0.01 M PBS/0.4% Triton solution. The sections were washed again in 0.01 M PBS over 30 min, incubated with 1:200 avidin-biotin-horseradish peroxidase

complex (Vectastain ABC Kit, Vector, Burlingame, CA) in 0.01 M PBS for 1 h, and then washed again in 0.01 M PBS over 30 min. The immune conjugate was visualized by placing the sections in 0.05% 3,3'-diaminobenzidine tetrahydrochloride (DAB) and 0.01% hydrogen peroxide in 0.01 M PBS. The visualization reaction was stopped by transferring the sections rapidly into cold PBS and washing thoroughly in 0.01 M PBS, after which the sections were mounted on gelatinized slides, dried, dehydrated through an ascending series of aqueous ethanol solutions (0%, 70%, 95%, 100% v/v), delipidated in HistoClear, and coverslipped with DePeX.

Neuronal loss was assessed as the absence of immunoreactive cells.

DEFINITION OF REINFORCEMENT SCHEDULES

The following reinforcement schedules are used or discussed in this thesis:

Continuous reinforcement (FR1). Every response is reinforced.

Extinction (ext). No response is reinforced.

Fixed ratio (FR). In an FR5 schedule, every fifth response is reinforced.

Variable ratio (VR). A ratio schedule in which the number of responses per reinforcer varies. If a single parameter is given (VR 5), the parameter is the mean number of responses required, with the probability distribution function (p.d.f.) unspecified.

Random ratio (RR). A random ratio schedule reinforces each response with a fixed probability. The parameter is the mean number of responses per reinforcement; thus an RR2 schedule is programmed as $P(\text{reinforcement} | \text{response}) = 0.5$.

Fixed time (FT). Reinforcement is delivered noncontingently at fixed, regular intervals.

Fixed interval (FI). After a successful response, reinforcement is not delivered until a certain time has elapsed. After this, the first response is reinforced.

Variable time (VT). A noncontingent schedule that delivers reinforcement after a certain time has elapsed since the previous reinforcement. If the interval is specified as a range (VT 10–30 s), a flat probability distribution of intervals is assumed (in this example, mean interval 20 s). If the interval is specified as a single number (VT 20 s) then this is the mean interval, with the p.d.f. unspecified.

Variable interval (VI). A contingent schedule; it is identical to a VT schedule, except that reinforcement is delivered for the first response following the interval.

Random time (RT). A form of VT schedule. Its objective is that the maximum interval between reinforcements is not specified (so that the probability of reinforcement does not increase to 1 at the end of the interval). A random time x schedule is typically programmed as $P(\text{reinforcement in each second}) = 1/x$. The probability of reinforcement is therefore independent of the time since last reinforcement. The expected number of reinforcements follows a binomial distribution, i.e. $P(k \text{ reinforcements in } n \text{ seconds}) = C(n, k)p^k(1-p)^{n-k}$, where $C(n, k) = n!/\{k!(n-k)!\}$. The mean number of reinforcements in n seconds is np and therefore the mean inter-reinforcement time in seconds is $n/np = 1/p$, the time for which the schedule is named. (The distribution of intervals follows the distribution function $P(\text{interval} = n) = (1-p)^{n-1}p$; this distribution declines continuously so the mode is 1 s; for $p = 1/30$ the median is 21 s and the mean is 30 s.)

Random interval (RI). A random interval schedule sets up reinforcement on an RT schedule; the next response is then reinforced.

BEHAVIOURAL APPARATUS

Unless otherwise specified, behavioural testing was conducted in eight identical operant chambers (30 × 24 × 30 cm; Med Instruments Inc, Georgia, Vermont, USA; Modular Test Cage model ENV-007CT). The chamber layout is shown in Figure 13. Each chamber was fitted with a 2.8 W overhead house light and two retractable levers, 16 cm apart and 7 cm above the grid floor, with a 2.8 W stimulus light (Ø 2.5 cm) above each lever and one located centrally (all 15 cm above the floor). The levers measured 4.5 cm (W) × 1.5 cm (D) and required a force of approximately 0.3 N to operate. In between the two levers was an alcove fitted with a 2.8 W lightbulb ('traylight', replaced

in some experiments by a 60 mcd diffused green LED, RS Components Ltd, UK), an infrared photodiode, a dipper that delivered 0.04 ml when elevated through a hole in the magazine floor, and a tray into which could be delivered food pellets. The chambers were enclosed within sound-attenuating boxes fitted with fans to provide air circulation. The apparatus was controlled by software written by RNC in Arachnid (Paul Fray Ltd, Cambridge), a real-time extension to BBC BASIC V running on an Acorn Archimedes series computer.

DATA ANALYSIS

Data collected by the chamber control programs were imported into a relational database (Microsoft Access 97) for case selection and analysed with SPSS 8.01, using principles based on Howell (1997). Graphical output was provided by Microsoft Excel 97 and SigmaPlot 5.0. All graphs show group means and error bars are ± 1 SEM unless otherwise stated.

Transformations. Skewed data, which violate the distribution requirement of analysis of variance, were subjected to appropriate transformations (Howell, 1997, §11.9). Count data (lever presses and locomotor activity counts), for which variance increases with the mean, were subjected to a square-root transformation. Homogeneity of variance was verified using Levene's test.

Analysis of variance. Behavioural data were subjected to analysis of variance (ANOVA) using a general linear model, using SPSS's Type III sum-of-squares method. Missing values were not estimated but excluded from analysis. All tests of significance were performed at $\alpha = .05$; full factorial models were used unless otherwise stated. ANOVA models are described using a form of Keppel's (1982) notation; that is, *dependent variable* = $A_2 \times (B_5 \times S)$ where A is a between-subjects factor with two levels and B is a within-subjects factor with five levels; S denotes subjects.

For repeated measures analyses, Mauchly's test of sphericity of the covariance matrix was applied and the degrees of freedom corrected to more conservative values using the Huynh-Feldt epsilon $\tilde{\epsilon}$ (Huynh & Feldt, 1970) for any terms involving factors in which the sphericity assumption was violated. Thus, the same analysis with and without sphericity correction would be reported as follows:

$$\begin{aligned} \text{Uncorrected: } F_{10,160} &= 2.047, p = .032 \\ \text{Corrected: } F_{4.83,77.3} &= 2.047, \tilde{\epsilon} = 0.483, p = .084 \end{aligned}$$

Post-hoc tests. Significant main effects of interest were investigated using pairwise comparisons with a Sidak correction. This is based on the observation that $\alpha_{\text{familywise}} = 1 - (1 - \alpha_{\text{each test}})^n$ when n tests are performed; the correction was made such that $\alpha_{\text{familywise}} = .05$.

Where main effects were found for between-subjects factors with three or more levels, post hoc comparisons were performed with the REGWQ range test (familywise $\alpha = 0.05$), or Dunnett's test in situations where several experimental treatments were compared with a single control group. These tests do not require the overall F for groups to be significant as they control the familywise error rate independently and test different hypotheses from the overall ANOVA, with different power (Howell, 1997, p. 351).

Where significant interactions were found following factorial analysis of variance, simple effects of a priori interest were calculated by one-way ANOVA and tested by hand against the pooled error term ($F = MS_{\text{factor}}/MS_{\text{pooled error}}$; critical values of F based on df_{factor} and $df_{\text{pooled error}}$). Multiple comparisons for simple effects were performed as described above but using the pooled error term.

Where significant interactions were found following repeated measures analysis, a pooled error term was used to test between-subjects simple effects of a priori interest, but separate error terms (i.e. plain one-way ANOVA) were used for within-subjects factors as sphericity corrections are inadequate if a pooled error term is used (Howell, 1997, p. 468).






Figure 13. **A:** Sketch of the operant chambers used in several tasks. **B:** Photograph of such an operant chamber. **C:** Close-up of a rat in the chamber.



Chapter 3.

Role of the anterior cingulate cortex in the control over behaviour by Pavlovian conditioned stimuli

Abstract. The anterior cingulate cortex (ACC) has been clearly implicated in stimulus–reward learning, but the exact contribution it makes to this process is not well understood. To address this issue, rats with lesions of peri- and postgenual ACC were tested using a variety of tasks to which stimulus–reward learning was expected to contribute. Unexpectedly, rats with ACC lesions learned to approach a food alcove during a stimulus that predicted imminent food delivery (temporally discriminated approach task), and subsequently responded normally for that stimulus in a test of conditioned reinforcement. They also exhibited normal conditioned freezing to an aversive CS that predicted footshock. Yet the same animals were impaired at autoshaping, a deficit observed before in ACC-lesioned animals. Furthermore, an autoshaping deficit was demonstrated when subjects received the lesion *after* training, though some behavioural recovery occurred. Additionally, the phenomenon of simple Pavlovian–instrumental transfer was intact following ACC lesions. In order to resolve the apparent discrepancy between the autoshaping deficit and the lack of a deficit on the temporally discriminated approach task, a new task was developed in which the approach behaviour was identical to that measured during the temporally discriminated approach task, but was under the control of two stimuli, only one of which was followed by reward. ACC-lesioned rats were impaired at the discrimination, approaching during both stimuli. It is suggested that this region of the ACC is not critical for stimulus–reward learning *per se*, but is required when multiple stimuli must be discriminated on the basis of their association with reward. Analogies with primate ACC are discussed.

INTRODUCTION

Delineation and connections of the anterior cingulate cortex (ACC) in the rat

The ACC is one of the three divisions of prefrontal cortex in the rat, the others being the agranular insular and orbitofrontal areas (Zilles & Wree, 1995). Definitions of this region vary. For example, Zilles & Wree (1995) define the ACC as comprising cortical subregions Cg1, Cg2, and Cg3, while Paxinos (1998) refers to Cg3 as prelimbic cortex (PrL). Previous lesion studies from this laboratory have used a definition based on vertical strips of cortex (Bussey *et al.*, 1996; 1997a; 1997b), discussed in detail by Bussey (1997b, p. 920). Figure 14 and Table 9 show these regions and various definitions of the ACC for comparison; the definitions of Bussey (1997b) will be followed in this thesis, except that Bussey’s ‘medial frontal cortex’ will be referred to as medial prefrontal cortex (mPFC).

These cortical regions have an extensive array of connections, summarized in Table 10. The most prominent efferent connections of the mPFC, ACC, and posterior cingulate cortex (PCC) are summarized by Bussey *et al.* (Bussey, 1996; Bussey *et al.*, 1997b) as follows. The mPFC, including anterior Cg1 and PrL, projects to the nucleus accumbens (Acb), mediodorsal nucleus of the thalamus (specifically the me-

dial part thereof: MDM), and the amygdala. The ACC (postgenual Cg1 and Cg2) projects to mediodorsal caudate, the lateral part of the mediodorsal nucleus of the thalamus (MDL), and the amygdala, while PCC projects to anteroventral and anterodorsal thalamic nuclei (AV, AD), the hippocampal formation (subiculum and parahippocampal cortex), visual cortex, and dorsal and mediodorsal striatum (for references, see Bussey *et al.*, 1997b).

Table 9. Definitions of the cingulate cortical divisions vary. See Figure 14.

Term	Definition of Zilles & Wree (1995)	Definition of Bussey <i>et al.</i> (1996; 1997a; 1997b), in terms of areas defined by Zilles (1985) and Zilles & Wree (1995)
medial prefrontal cortex	(this term is used descriptively to include infralimbic cortex, Zilles & Wree, 1995, p. 653, but is not defined)	Cg3; Cg1 rostral to genu of the corpus callosum. This is equivalent to PrL plus rostral Cg1 in the atlas of Paxinos & Watson (1998); see Figure 14.
anterior cingulate cortex	Cg1–3	Cg1 and Cg2 caudal to the genu (dorsal and ventral ACC respectively)
posterior cingulate cortex	synonym for retrosplenial cortex	RSA and RSG rostral to the splenium of the corpus callosum
retrosplenial cortex	granular and agranular retrosplenial cortex (RSG, RSA)	RSA and RSG caudal to the splenium of the corpus callosum


Table 10. Connections of the anterior cingulate cortex. From Zilles (1995, p. 654); additional data (*) from Brog *et al.* (1993); abbreviation key also from Price (1995) and Paxinos & Watson (1998).

Area	Afferent input	Efferent output
All areas (Cg1–3)	basal nucleus of Meynert (B); basolateral amygdala (BL); caudal interstitial nucleus of the medial longitudinal fasciculus (CI); dorsal raphé nucleus (DR); median raphé nucleus (MnR); intralaminar thalamic nuclei; locus coeruleus (LC); lateral hypothalamic area (LH); mediodorsal thalamic nucleus (MD); parabrachial pigmented nucleus (PBP); substantia nigra (SN); ventromedial thalamic nucleus (VM); ventral tegmental area (VTA); zona incerta (ZI); centrolateral thalamic nucleus (CL); laterodorsal thalamic nucleus (LDs); periventricular hypothalamic nucleus (Pe); infralimbic cortex (IL); cingulate areas Cg1–3 contralaterally; agranular insular cortical areas	caudal interstitial nucleus of the medial longitudinal fasciculus (CI); intralaminar thalamic nuclei; lateral habenular nucleus (LHb); pontine nuclei (Pn); anterior pretectal nucleus (APT); mediodorsal thalamic nucleus (MD); periventricular hypothalamic nucleus (Pe); median raphé nucleus (MnR); reticular thalamic nucleus (Rt); ventromedial thalamic nucleus (VM); superior colliculus (SC); agranular insular cortex, ventral part (AIV); entorhinal cortex (Ent); presubiculum (PrS); Cg1–3 contralaterally nucleus accumbens core (AcbC)*; nucleus accumbens shell (AcbSh)*
Cg1–2	anteromedial thalamic nucleus (AM); lateral posterior thalamic nucleus (LP)	anteromedial thalamic nucleus (AM); caudate-putamen (CPu); lateral dorsal thalamic nucleus (LD); retrosplenial granular cortex (RSG); retrosplenial agranular cortex (RSA)
Cg1	gigantocellular reticular nucleus, ventral part (GiV); lateral paragigantocellular nucleus (LPG); lateral reticular nucleus (LRt); medial occipital area 2 (Oc2M); occipital cortex, area 1 (Oc1); retrosplenial granular cortex (RSG); retrosplenial agranular cortex (RSA)	agranular insular cortex, dorsal part (AID); agranular insular cortex, ventral part (AIV); pontine reticular nucleus (PnC, PnQ)
Cg2	medial occipital area 2 (Oc2M); retrosplenial granular cortex (RSG); retrosplenial agranular cortex (RSA)	parietal area 2 (Par2); cingulate areas Cg1 and Cg3; retrosplenial granular cortex (RSG)
Cg3 (PrL)	paratenial thalamic nucleus (PT)	amygdaloid nuclei; lateral hypothalamic area (LH); midline thalamic nuclei; paratenial thalamic nucleus (PT); substantia nigra (SN); mesencephalic tegmentum; nucleus of the solitary tract (Sol); olfactory tubercle (Tu); ventral tegmental area (VTA); agranular insular cortical areas; perirhinal cortex (PRh); piriform cortex (Pir); substantia innominata (SI); nucleus of the horizontal limb of the diagonal band (HDB)


In addition to its reciprocal connections with other areas of prefrontal cortex and the basolateral amygdala, the ACC has both direct and indirect connections to the ventral striatum (see Alexander *et al.*, 1986). Not only does the ACC project to the mediodorsal striatum (Zilles & Wree, 1995, p. 654), but both anterior cingulate and prelimbic cortex project to the core and rostral pole of the Acb (McGeorge & Faull, 1989; Zahm & Brog, 1992; Brog *et al.*, 1993; Parkinson, 1998) (see also Heimer *et al.*, 1995, pp. 600–601). The ACC also receives major dopaminergic input from the VTA (Fallon & Loughlin, 1995). Not only does the ACC provide a major input to the ventral striatum but this ‘limbic loop’ of the basal ganglia

projects via the ventral pallidum back to the ACC as well as the mPFC (Alexander *et al.*, 1986). This is the basis of an anatomical argument that the mPFC and ACC are the primary cortical structures whose information content is affected by the Acb (see Heimer *et al.*, 1995, p. 613).

A



B



C





Figure 14. Midline cortex in the rat, including the cingulate areas. **A:** Figure taken from Paxinos (1998); sagittal section at 0.4 mm lateral to the midline. Anterior is to the left. (Selected abbreviations: Cg1/Cg2, cingulate cortex areas 1/2; PrL, prelimbic cortex; IL, infralimbic cortex; M2, secondary motor cortex; FrA, frontal association cortex; MO, medial orbital cortex; cc, corpus callosum; gcc, genu of the corpus callosum; scc, splenium of the corpus callosum.) **B:** Figure taken from Zilles & Wree (1995), p. 650. Unshaded areas, including Cg1, are neocortex; shaded areas, including Cg2–3, are periallo- and proisocortex. (VO, ventral orbital area; OB, olfactory bulb.) **C:** Figure taken from Bussey (1996), indicating the area of the lesions used by Bussey *et al.* (1996; 1997a; 1997b). Vertical strips, from anterior to posterior, represent medial frontal cortex (medial prefrontal cortex, mPFC), anterior cingulate cortex (ACC) and posterior cingulate cortex (PCC). For comparison with other lesion studies of the rat, note that Weissenborn *et al.* (1997) used the same coordinates as Bussey *et al.* for their post-genual ACC lesions, except that the most anterior injection of toxin was 0.1 mm more caudal. Muir *et al.* (1996) used a different ear bar setting but also aimed at post-genual ACC. **D:** Figure indicating the region of the ACC targeted in the present study, encompassing the perigenual area.

D



Involvement of the rat ACC in stimulus-reinforcer association

A range of studies have implicated the ACC in stimulus–reinforcer association, using both appetitive and aversive tasks. The ACC receives nociceptive information and is involved in the coordination of autonomic responses (Neafsey *et al.*, 1993; Fisk & Wyss, 1997; Hsu & Shyu, 1997); early studies found that aspirative lesions of the ACC attenuated classically conditioned bradycardia in the rabbit (Buchanan & Powell, 1982a). In the rabbit, the ACC is also involved in active avoidance behaviour. Using a task in which rabbits must learn to step in response to a tone CS+ in order to avoid a shock, while ignoring a different tone (CS–), Gabriel *et al.* have shown electrophysiologically that discriminated neuronal activity (discharge to the CS+ but not the CS–) occurs early in avoidance training (Gabriel *et al.*, 1980a; Gabriel *et al.*, 1980b; Gabriel & Orona, 1982; Gabriel *et al.*, 1991b). Lesions of the ACC impair the avoidance response (Gabriel *et al.*, 1991a; Gabriel, 1993), attributed to the loss of associative information about the significance of a discrete CS (Gabriel *et al.*, 1980a, pp. 158–163/219–221).

In the rat, the ACC has more often been studied using appetitive tasks, which also suggest that it has a role in stimulus–reinforcer association. For example, Bussey *et al.* (1997b) found that lesions of the ACC impaired the acquisition of an eight-pair concurrent discrimination task, in which subjects must learn which stimulus in each of eight pairs of complex visual stimuli must be selected in order to obtain reward. Additionally, Bussey *et al.* have reported that ACC lesions *facilitate* early learning in a conditional visual discrimination (CVD) task (Bussey *et al.*, 1996), though not in all circumstances (Bussey *et al.*, 1997b). In this task, subjects must respond in one way to stimulus A, and in another way to stimulus B; the reward is identical in both situations. This task cannot be solved by the formation of stimulus–reinforcer associations, but is soluble through stimulus–response association. Bussey *et al.* (1996) have suggested that the facilitation they observed with ACC lesions was due to the loss of a stimulus–reward system that normally competes with a stimulus–response system in the PCC during learning or behavioural expression (Bussey *et al.*, 1996; 1997b).

Few of the tasks described so far directly address the question of whether the ACC is involved in classical conditioning. To examine classical conditioning in isolation, it is necessary either to ensure that the animal's behaviour is uncorrelated with the presentation or receipt of the reinforcer, or that the instrumental behaviour that produces the reinforcer is directly opposed to the classically conditioned response elicited by the CS (omission schedules; Sheffield, 1965). Autoshaping, in which animals approach a stimulus that predicts reward, is a relatively selective test of Pavlovian learning. Autoshaping was originally demonstrated in pigeons by Brown & Jenkins (1968), who illuminated a response key and delivered food immediately afterwards. Regardless of the fact that responding had no effect on food delivery, the subjects reliably approached and pecked the key. There is no instrumental contingency specified in the task, and as the autoshaped response is to the stimulus rather than the place of food delivery there is little opportunity for 'implicit' instrumental response–reward associations. Furthermore, the nature of the autoshaped response is specific to the reinforcer (Jenkins & Moore, 1973) and subjects will immediately approach the CS+ following training in which approach has been prevented by a barrier (Browne, 1976). Finally, alteration of the contingencies so that approach prevents reward delivery — an omission schedule — fails to eliminate responding to the CS+ (Williams & Williams, 1969). In the study of Bussey *et al.* (1997a), not only did control rats fail to alter their responding when an omission contingency was introduced, but the ratio of CS+/CS– approaches *increased* as overall responding extinguished. Thus there is strong evidence that normal animals' behaviour is governed by Pavlovian associations in this procedure.

Bussey *et al.* (1997a) found that lesions of the ACC significantly impaired the acquisition of an autoshaping task. In their task, a visual stimulus (CS+) is presented on a computer screen and followed by

delivery of food at a different location. A second stimulus (CS⁻) is also presented, but never followed by food. Though the subject's behaviour has no effect on food delivery, normal rats develop a conditioned response in which they selectively approach the CS predictive of food before returning to the food hopper to retrieve the primary reward. In contrast, rats with lesions of the ACC fail to discriminate, approaching the CS⁺ and CS⁻ equally. It is intriguing to note, however, that the lack of discrimination in ACC-lesioned rats takes the form of increased responding to the CS⁻, rather than decreased responding to the CS⁺ (Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c). As ACC-lesioned rats have been shown to be somewhat 'disinhibited', reflected in their tendency to make inappropriate premature responses in a test of sustained attention (Muir *et al.*, 1996), it is unclear whether their impairment in the autoshaping task is due to a failure to learn CS-US associations entirely (coupled with a tendency to over-respond to both the CS⁺ and the CS⁻) or a specific failure to inhibit responding to unrewarded stimuli. In fact, it is presently unknown whether the autoshaping impairment represents failure to learn at all, or simply to express learning that occurs in other brain regions. It seems unlikely, however, that the deficit is attentional, as ACC lesions do not impair the accuracy of visual attentional function (Muir *et al.*, 1996).

The ACC projects to the Acb; this projection, and the Acb itself, is also critical for the development of autoshaping, suggesting that information stored in or retrieved by the ACC gains access to locomotor response systems via the Acb (Parkinson *et al.*, 1996; Parkinson *et al.*, 2000c). In addition, the Acb is involved in another aspect of Pavlovian conditioning: conditioned reinforcement. Following the discovery that intra-accumbens injection of the psychostimulant *d*-amphetamine selectively enhances responding for conditioned reinforcement in a dose-dependent manner (Taylor & Robbins, 1984), attention has focused on the neural structures that convey information regarding the value of conditioned reinforcers to the Acb. The major cortical inputs to Acb are the basolateral amygdala (BLA), the entorhinal cortex and hippocampus (largely via the ventral subiculum), the mPFC (including prelimbic cortex, Cg3), and the ACC (Cg1-2) (Zahm & Brog, 1992; Brog *et al.*, 1993; Parkinson, 1998). Lesions of the ventral subiculum and mPFC do not impair responding for conditioned reinforcement (Burns *et al.*, 1993), but lesions of the BLA do so dramatically (Cador *et al.*, 1989; Burns *et al.*, 1993). It is not presently known whether the ACC is also required for conditioned reinforcement. Given that the ACC projects both to the BLA and the Acb, and has been implicated in stimulus-reward association, it is clearly of interest to establish whether it plays a role in the ability of neutral stimuli to gain conditioned reinforcing properties.

In order to address these questions, the present study investigated the effects of excitotoxic lesions of the ACC on the acquisition of a simple, temporally discriminated approach task. In this task, a single stimulus predicted the delivery of food at the same location. Following establishment of this stimulus as an appetitive CS, the subjects were allowed to respond for the same stimulus in the absence of any primary reward, the CS now acting as a conditioned reinforcer. At the same time, the effects of intra-accumbens amphetamine injections were examined in control and ACC-lesioned subjects; in addition to promoting responding in extinction (Robbins, 1976), this technique allowed the establishment of the amphetamine dose-response curve for comparison with previous lesion studies. Although the ability of a stimulus to act as a conditioned reinforcer indicates that it has entered into a Pavlovian association with its US (see Mackintosh, 1983, p. 15), the temporally discriminated approach task used to establish this association was not a pure measure of Pavlovian conditioning. Though the CS predicted the arrival of food, allowing approach behaviour to be classically conditioned to the CS, the CS might also have served as a discriminative stimulus (S^D), signalling that an instrumental contingency existed between approach behaviour and food acquisition. Therefore, the effects of ACC lesions were also tested using a number of purer measures of Pavlovian conditioning: autoshaping, conditioned freezing (a measure of aversive con-

ditioning), and the phenomenon of Pavlovian–instrumental transfer (Estes, 1948; Lovibond, 1983), in which a Pavlovian CS enhances ongoing instrumental responding. As it is unclear whether the autoshaping deficit previously reported in ACC-lesioned rats (Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c) represents a failure of learning or of performance (see Table 6, p. 36), lesions were also made after the acquisition of autoshaped behaviour, to test whether the ACC is required for performance in well-trained animals.

Lesion methods and sites within the ACC

In order to evaluate the function of the cingulate cortex by means of lesion studies, axon-sparing (excitotoxic) lesions must be used, as damage to the underlying cingulum bundle can itself produce significant behavioural impairments (Meunier & Destrade, 1988; Warburton *et al.*, 1998). All experiments reported here use the excitotoxic technique.

A pilot study using the same ACC lesion coordinates as Bussey *et al.* (1997a) revealed a non-significant trend towards an impairment in discriminated approach behaviour very early in training (sham $n = 11$, lesion $n = 9$), and a trend towards an impairment in responding for conditioned reinforcement in those animals whose lesions extended anteriorly to the perigenual region of the ACC (sham $n = 10$, lesion subgroup $n = 5$). This same subgroup demonstrated the poorest autoshaping, and the projections from the ACC to the AcbC are known to arise from the perigenual region (McGeorge & Faull, 1989; Brog *et al.*, 1993; Parkinson, 1998). Therefore, all experiments in this thesis used lesions of the ACC centred on the perigenual region.

EXPERIMENT 1: EFFECTS OF ACC LESIONS ON TEMPORALLY DISCRIMINATED APPROACH, RESPONDING FOR CONDITIONED REINFORCEMENT, AND FEAR CONDITIONING TO A DISCRETE CUE

Methods

Overview

Twenty-two male hooded Lister rats received lesions of perigenual ACC (group ACCX, $n = 12$) or sham lesions (group sham, $n = 10$), with all animals additionally receiving bilateral cannulae aimed at the Acb. They weighed 295–390 g at the time of surgery. Following recovery, they were maintained at 85% of their free-feeding mass and underwent the following behavioural procedures, in order: (1) temporally discriminated approach to a stimulus predictive of sucrose; (2) acquisition of a new response with conditioned reinforcement, with intra-accumbens amphetamine injections; (3) autoshaping; (4) a sucrose consumption test in the home cages; (5) locomotor activity testing in a novel environment; (6) acquisition of freezing to a stimulus predictive of footshock. During the conditioned freezing test they were allowed free access to food. After this they were killed and perfused for histology.

Housing conditions, operative techniques and stereotaxic coordinates are described fully in the *Methods* chapter.

Temporally discriminated approach

Four operant chambers were used for the acquisition of discriminated approach and instrumental responding phases; for this task they were fitted with a 2.8 W bulb traylight and the pellet tray was not present.

No levers were extended during this task. At the start of any session, the houselight was on, the traylight was off and the dipper was not raised. This phase lasted for a variable interval (VI) of 30–90 seconds, randomly chosen for each cycle of CS–US presentation. This was followed by a CS: the houselight was switched off and the traylight was switched on for a period of 5 s. The CS was immediately followed by the US: the traylight was switched off, the houselight was switched back on, and the dipper was raised for 5 s to deliver 10% w/v sucrose solution. The dipper was then lowered to return the chamber to the starting state and the next VI began.

Animals were trained for 11 sessions with one session per day. In each session, the subjects received 30 presentations of the CS and US. For each period (VI, CS, US), the number of entries into the food alcove and the time spent in the alcove were recorded. The proportions of the CS and VI periods that the subject spent in the alcove were combined to calculate an approach ratio equal to $(\text{CSproportion} \div (\text{CSproportion} + \text{VIproportion}))$, used as a measure of conditioning to the CS.

Acquisition of a new response with conditioned reinforcement

This task was conducted in the same apparatus as the temporally discriminated approach task. Test sessions were conducted in extinction, and immediately followed bilateral administration of one of 4 doses of intra-accumbens D-amphetamine sulphate (Sigma, UK; 0, 3, 10 and 20 μg in 1 μl of 0.1 M sterile phosphate buffer, pH 7.4). Doses were counterbalanced in a Latin square design to eliminate differential carryover effects and separated by 24 h. The Latin square was of a digram-balanced design (Keppel, 1991, p. 339), in which each condition immediately precedes and follows the other conditions once (e.g. 1234, 3142, 2413, 4321). Sensitization to amphetamine does not occur with repeated administration into the Acb (Cador *et al.*, 1995), so further spacing of doses was not required.

A session began when the subject nose-poked in the central alcove, and lasted 30 minutes. Initially, the houselight was switched on, the traylight was off, and both levers were extended. Responding on one of the levers, the CRf lever, resulted in the presentation of an abbreviated version of the previous conditioned stimulus with a probability of 0.5 (a random ratio 2 schedule). To produce this stimulus, the houselight was switched off and the traylight was switched on for 0.5 s, after which the lights were returned to the initial state and the empty dipper was raised for 0.3 s; this stimulus is known to function well as a conditioned reinforcer (Burns *et al.*, 1993). Responding on the other (NCRf) lever had no programmed consequence. The lever assignment (left/right) was counterbalanced across rats.

Alcove approach frequency and duration were recorded, together with all lever-pressing activity. All measures of behaviour were recorded in six 5-minute bins.

Intracranial infusion during conditioned reinforcement test

Before the first test day, all rats were given a preliminary infusion of vehicle and returned to the home cage to familiarize them with the hand-held infusion procedure and to minimize non-specific effects of inserting the infusion cannulae during subsequent test sessions. On the first infusion only, these effects are noticeable; many animals become slightly agitated near the end of the infusion period and a few react briefly as the injector is removed. This has been observed with a variety of intracranial cannula sites (F. Passetti, personal communication, 1998).

Intra-accumbens infusions were performed by inserting two 28-gauge infusion cannulae (\varnothing 0.36 mm external, 0.18 mm internal; model C313I, Plastics One, Roanoke, Illinois, USA; supplied by Semat Technical Ltd, St Albans, UK) through the chronically implanted 22-gauge guide cannulae of gently hand-restrained subjects. The infusion cannulae were 15.0 mm long so as to allow them to protrude 2.0 mm beyond the tips of the guide cannulae; they were connected by polyethylene (PE50) tubing to two 5- μ l syringes (SGE Ltd, Milton Keynes, UK) mounted on a Harvard Apparatus (Edenbridge, UK) infusion pump. Amphetamine was infused in a volume of 1 μ l per side over a 2-minute period. After this, 2 minutes were allowed for diffusion away from the site of the cannulae to occur, before the cannulae were removed and replaced by occluders and behavioural testing began. Animals were held during the infusion but otherwise allowed to move freely.

Autoshaping

Apparatus. Autoshaping was assessed in the apparatus shown in Figure 15 and is described fully in Bussey *et al.* (1997a). Briefly, the apparatus consists of a 48 \times 30 \times 30 cm testing chamber with a display screen on one wall and a pellet dispenser located centrally in front of the display. Pressure-sensitive areas of floor (each 14 \times 10 cm) were located directly in front of the display, to the left and right of the dispenser, and also centrally at the rear of the chamber. The apparatus was controlled by software written in BBC BASIC by T.J. Bussey, running on a BBC Master series computer.

Pretraining. Rats were first given one session in order to habituate to the test chamber and to collect 45-mg food pellets (Rodent Diet Formula P, Noyes, Lancaster, NH) from the food receptacle. The houselight was illuminated and subjects were placed in the chamber for 5 min with 4–5 pellets placed in and around the dispenser. After this, pellets were delivered on a VT 0–40 s schedule for 15 min.

Acquisition (CS⁺→food, CS⁻→0). On the next day, rats were trained to associate stimuli with the delivery of pellets. Stimuli consisted of 8 \times 18 cm white vertical rectangles displayed on the left and right of the screen for 10 s. One was designated the CS⁺ and the other the CS⁻, counterbalanced between subjects. A trial consisted of presentation of both the CS⁺ and CS⁻ in a randomized order. Following a VI of 10–40s, the program waited for the rat to be located centrally at the rear of the chamber; this eliminated chance approach to the stimuli, ensured equal stimulus sampling and allowed accurate measurement of approach latency. One stimulus was then presented for 10 s. The CS⁺ was always followed immediately by the delivery of food; the CS⁻ was never followed by food. After this, another VI followed, the program waited for the rat to return to the rear of the chamber, and the other stimulus was presented. This procedure ensured that the minimum time between CS⁺ and CS⁻ presentation was 10 s, and that there were never more than two consecutive presentations of either the CS⁺ or the CS⁻.

When a stimulus was presented, activation of one of the two floor panels in front of the screen was scored as an approach, and no further approaches were scored during that stimulus presentation. The rat may therefore make four kinds of active response: approach to the CS⁺, approach to the CS⁻, approach to the location of the CS⁺ during CS⁻ presentation, and approach to the location of the CS⁻ during CS⁺ presentation. Rats were trained for a total of 100 trials (two days with 50 trials per day). Approaches to the CS⁺ and the CS⁻ were scored in blocks of 10 trials and mean approach latency was calculated over 100 trials (Bussey *et al.*, 1997a). (In some previous studies using this task, trials on which the approach latency was under 10 cs were excluded as representing equipment failure. Covert observation revealed that such latencies were genuinely attainable, because the software took a perceptible fraction

of a second to draw each stimulus, and began timing at the point when drawing was complete; if, while the stimulus was being drawn, the rat ran to the front of the chamber, very short latencies were reported. Therefore, these trials were included in the analysis.) Data were analysed as CS⁺/CS⁻ approach scores, as difference scores (CS⁺approaches – CS⁻approaches) (after Bussey *et al.*, 1997a) and as the ratio (CS⁺approaches) ÷ (CS⁺approaches + CS⁻approaches), a measure of stimulus discrimination that is relatively independent of absolute approach activity.




Figure 15. Autoshaping apparatus. (From Bussey *et al.*, 1997a.)

Probe trials (CS⁺ and CS⁻). After acquisition, a probe test was performed, consisting of 20 trials in which the CS⁺ and CS⁻ were presented simultaneously and approaches were measured. Food was not delivered, so this test constituted an extinction trial to the CS⁺, while the CS⁻ was still a perfect predictor of food absence. The probe test was intended to be a more sensitive test than the acquisition task (in which the subject might form CS–US associations perfectly and yet approach all stimuli), as it forced the subject to make a choice between the CS⁺ and the CS⁻.

Omission training. Finally, the contingencies were altered such that approaches to the CS⁺ prevented the delivery of a food pellet. This manipulation introduced an instrumental contingency directly opposed to the approach response. All other parameters remained the same as in the acquisition phase. There were 50 presentations of the CS⁺ and of the CS⁻ per session and two sessions were given. As before, only initial approaches were scored; ‘successful’ omission trials were those in which the CS⁺ was presented and the subject first approached the CS⁻, or failed to approach either stimulus. (In fact, the program incorrectly omitted reward even if the rat first approached the CS⁻ and later wandered over to the CS⁺ side while the stimulus was still present — and, I believe, if contact with the CS⁻ itself was made. However, these were extremely rare events.)

Sucrose consumption

In order to assess alterations in primary motivation, all animals were given a sucrose consumption test while food-deprived. Intake of 10% sucrose solution was measured during 1 h of free access in the home cages with a single subject present.

Locomotor activity in a novel environment

Locomotor activity was measured in wire mesh cages, 25 (W) × 40 (D) × 18 (H) cm, equipped with two horizontal photocell beams situated 1 cm from the floor that enabled movements along the long axis of the cage to be registered. Subjects were placed in these cages, which were initially unfamiliar to them, and their activity was recorded for 2 h. All animals were tested in the food-deprived state.

Fear conditioning to a discrete cue

Fear conditioning was carried out using two distinctive experimental contexts, termed Light and Dark. The Light context consisted of a 20 (W) × 21 (D) × 21 (H) cm chamber fitted with white and steel walls on three sides and a fourth transparent Perspex wall that also served as a door. The floor consisted of a steel grid (bars 0.75 cm apart) on top of which was placed a transparent Perspex sheet; under the grid was a tray of sawdust. There was a white 2.5-W houselight in the centre of the chamber's ceiling. In front of the transparent wall was a Sony VHS-C video camera on a tripod; the room was illuminated by a white fluorescent ceiling lamp at moderate intensity. The Dark context consisted of a 35 (W) × 25 (D) × 40 (H) cm chamber in a room illuminated only by a 40-W red incandescent lamp. The chamber had four black Perspex walls and a transparent ceiling; it had a red 2.5-W houselight and a steel grid floor (bars 1 cm apart), 3 cm above a steel tray scented with a small quantity of apricot-scented oil (Crabtree and Evelyn, UK). A shock scrambler (model 521C, Campden Instruments, Loughborough, UK) could deliver brief electric shock to the grid floor. Both contexts were equipped with identical 80-dB clicker relays.

Contexts were made more discriminable by ensuring a unique time of day was paired with each environment (counterbalanced across rats); for example, half of the rats only ever experienced the Light context in the morning and the Dark context in the afternoon.

On days 1–3 of the experiment, subjects were pre-exposed by being placed for 25 min in each context. On day 4, they were placed in the Dark context, where they received 5 presentations of a 10-s clicker CS (5 Hz cycle for a 10 Hz click rate) terminating in a shock of 0.5 mA lasting 0.5 s. The interval between presentations was 4 ± 1 min and the animals were in the context for 30 min. On day 5, subjects were placed in the Light context and their behaviour was videotaped. After 5 min of CS absence, the clicker CS was played continuously for 10 min. Freezing activity was assessed by an observer scoring the tapes in 5-s activity bins, using a stringent criterion: if and only if the animal was motionless apart from respiratory movements for the full 5 s, the bin was scored as 'freezing'. The calculated measure was the percentage of bins spent freezing; the 2 minutes preceding CS onset were compared with the 8 minutes following CS onset.

Results


One subject in the ACCX group (subject E2) lost its cannulae and was killed. There were 3 other postoperative deaths (E1, E7, E9). After histological analysis, all lesions were found to be complete, leaving 8 animals in the ACCX group (subjects E3, E4, E5, E6, E8, E10, E11, E12) and 10 in the sham group (subjects E13, E14, E15, E16, E17, E18, E19, E20, E21, E22), of which respectively 6 and 10 also had injection sites correctly located within the Acb (all but subjects E5 and E8). Data from all animals with valid lesions were analysed, except for the conditioned reinforcement test, for which only data from animals with valid lesions and valid cannulae placements were used.

Histology


In this group of ACC-lesioned subjects, neuronal loss and associated gliosis extended from ~2.5 mm anterior to bregma to ~0.3 mm posterior to bregma, destroying perigenual Cg1 and Cg2; there was minimal damage to PrL (a few subjects exhibited a small degree of neuronal loss in the most dorsal aspect of PrL). IL and PCC were undamaged, as was the corpus callosum. Photomicrographs of the ACC in a sham-operated and a lesioned rat are shown in Figure 16; this material was typical of lesions in this group. Schematics depicting the largest and smallest extent of the lesions are shown in Figure 17. Photomicrographs of the location of the intra-accumbens guide cannulae and injector tip locations are shown in Figure 18, indicating the minimum and maximum amount of damage done by the guide cannulae, while Figure 19 presents a schematic of the injector tip locations in the two groups.

Anterior cingulate cortex: photomicrographs

A. Sham (low power)



B. Sham (high power)



C. Lesion (low power)



D. Lesion (high power)




Figure 16. Lesion of the ACC: photomicrographs of coronal brain sections, approximately 0.5 mm anterior to bregma, stained with cresyl violet. **A & B:** sham-operated rat (cc, corpus callosum; LV, lateral ventricle; Cg1/Cg2, cingulate areas 1/2; M2, secondary motor cortex; M1, primary motor cortex). **C & D:** ACC-lesioned rat; dotted lines mark the borders of the lesion. **Left-hand** panels (A & C) are low-magnification view (scale bars are 1 mm); **right-hand** panels are high-magnification views (scale bars are 0.1 mm). Arrowheads mark identical structures in the respective low- and high-power views.

Schematic of lesions




Figure 17. Schematic of lesions of the ACC (subjects E3, E4, E5, E6, E8, E10, E11, E12). Grey shading indicates the extent of the largest area of neuronal loss, and black the smallest. Diagrams are taken from Paxinos & Watson (1998). The pair of vertical lines in the sagittal schematic (top right) indicate the anterior and posterior limits of the series of coronal schematics (main part of figure)

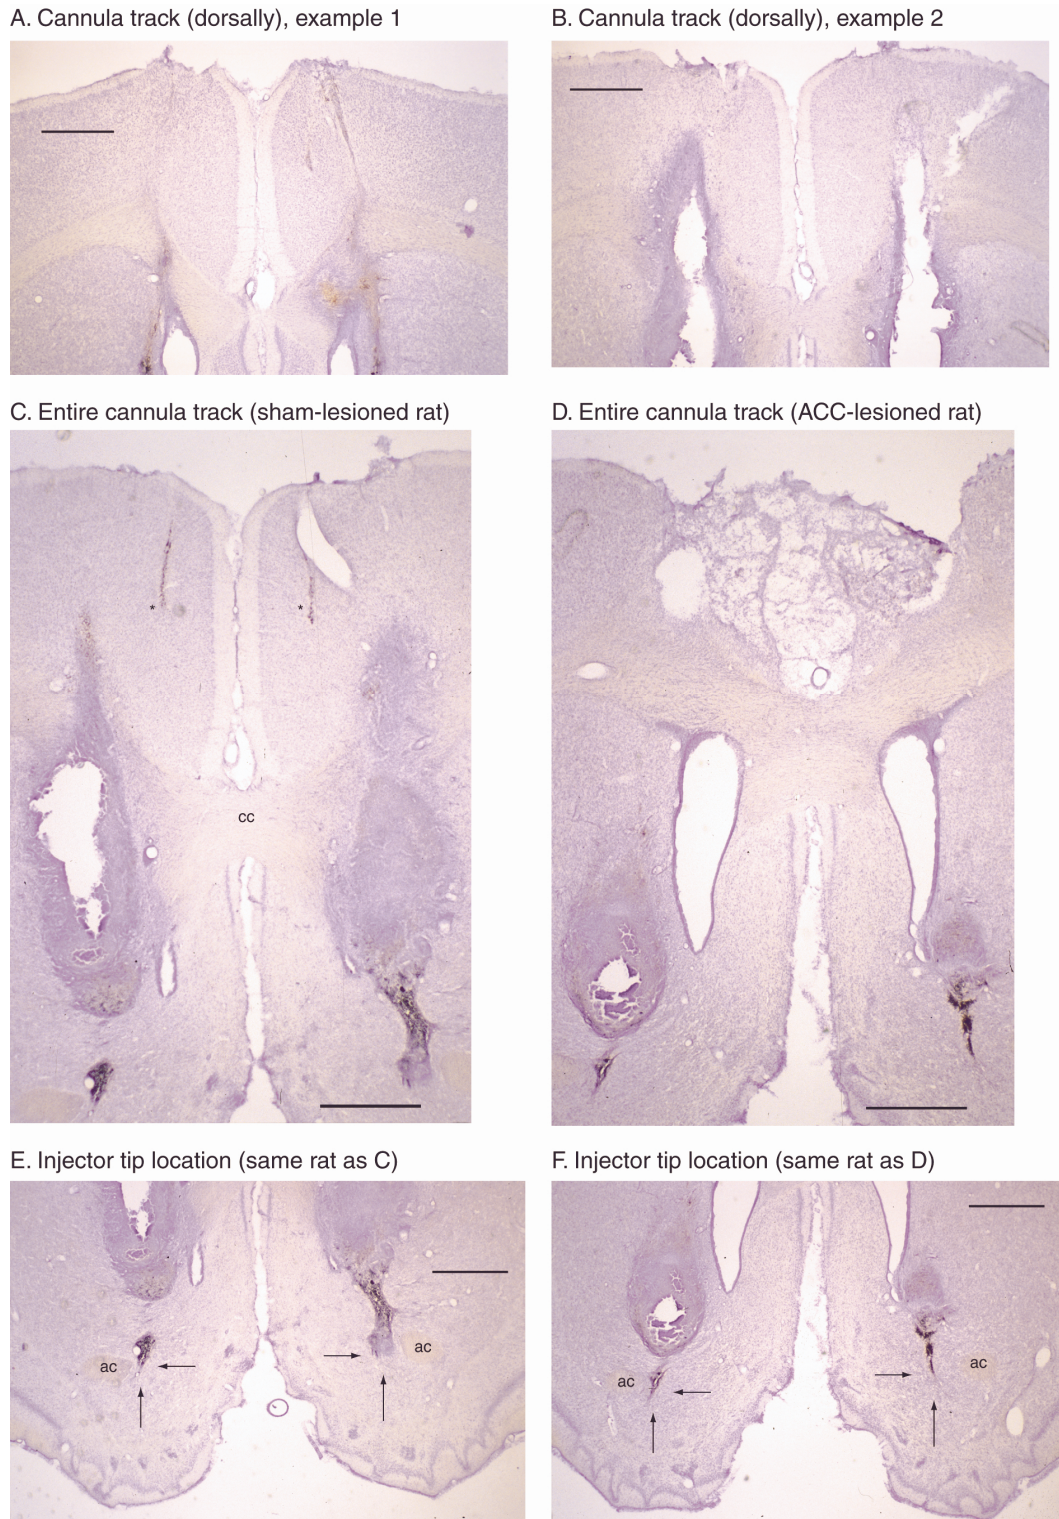
Nucleus accumbens cannulae: photomicrographs of cannula tracks

Figure 18. Location of intra-accumbens guide cannulae and injector tips. All sections are at approximately 1.6 mm anterior to bregma. **A:** Dorsal part of cannula tracks in a rat with minimal track damage. **B:** Dorsal track of cannulae in a rat with more pronounced track damage. **C:** View of cannula tracks and location of injector tips within the Acb, in a rat that received sham anterior cingulate surgery. The ACC is intact, and needle tracks are visible where the vehicle was injected (*) (cc, corpus callosum). **D:** Location of injector tips in a cingulate-lesioned rat. The excitotoxic lesion of the ACC is clearly visible (compare Figure 16). **E:** View of the injector tip location in the Acb (same rat as C). Perpendicular arrows point to the tip location in each hemisphere (ac, anterior commissure). **F:** Close-up of the tip location in the Acb (same rat as D). All scale bars are 1 mm.

Nucleus accumbens cannulae: schematic of cannula placements



Figure 19. Location of the tips of injection cannulae within the Acb. Red triangles indicate subjects with lesions of the ACC (subjects E3, E4, E6, E10, E11, E12). Black crosses indicate sham-operated control subjects (E13, E14, E15, E16, E17, E18, E19, E20, E21, E22). Diagrams are taken from the atlas of Paxinos & Watson (1998).

Temporally discriminated approach

All animals learned to approach the alcove during the CS selectively; the lesioned and sham groups did not differ in any respect, as shown in Figure 20. All dependent variables were analysed using the model group \times (session \times S). Analysis of the approach ratios revealed a main effect of session ($F_{6.887,110.187} = 92.821$, $\tilde{\epsilon} = .689$, $p < .001$), reflecting a selective increase in approach during the CS, but there was no effect of group ($F < 1$, NS) and no group \times session interaction ($F_{6.887,110.187} = 1.253$, $\tilde{\epsilon} = .689$, NS). A similar pattern was observed for the proportion of the CS spent nosepoking (session: $F_{6.781,108.493} = 42.108$, $\tilde{\epsilon} = .678$, $p < .001$; group: $F_{1,16} = 1.289$, NS; group \times session: $F < 1$, NS), for the percentage of trials on which the CS was approached at least once (session: $F_{10,160} = 76.876$, $p < .001$; group: $F < 1$, NS; group \times session: $F < 1$, NS) and for the time spent approaching the food alcove during the VI (session: $F_{6.043,96.686} = 6.562$, $\tilde{\epsilon} = .604$, $p < .001$; group: $F_{1,16} = 1.698$, NS; group \times session: $F < 1$, NS). It was clear that the learning resulted in dramatically improved access to the US (Figure 20E) and again there was no effect of the lesion on this measure (session: $F_{6.178,98.841} = 90.717$, $\tilde{\epsilon} = .618$, $p < .001$; group: $F < 1$, NS; group \times session: $F < 1$, NS).




Figure 20. Temporally discriminated approach behaviour was unaffected by lesions of the ACC. **A:** Approach ratio. This ratio is calculated as the proportion of the CS time spent nosepoking divided by the sum of the proportions of CS and VI time spent poking; this measure is therefore independent of CS and VI durations. A measure of 0.5 indicates that nosepoking was evenly distributed between the CS and VI, while a ratio of 1 indicates that responding occurred solely during the CS. **B:** Approach during the CS: the proportion of time spent nosepoking during CS presentation. **C:** Percentage of trials on which the CS was approached at least once. **D (overleaf):** Approach during the VI, as a proportion of VI duration. **E (overleaf):** Approach during the US, as a proportion of US duration.




Figure 20 (continued). See previous page for caption.

Responding for conditioned reinforcement

Animals responded more on the lever producing the conditioned reinforcer (CRf lever) than the control (NCRf) lever, and responding for the CRf was dose-dependently and selectively potentiated by intra-accumbens amphetamine, but lesioned and sham groups did not differ (Figure 21A). Lever-press data were subjected to a square-root transformation and analysed using the model group \times (lever \times dose \times S). Subjects responded more on the CRf than the NCRf lever (effect of lever, $F_{1,14} = 29.422$, $p < .001$). Amphetamine selectively potentiated responding on the CRf lever (lever \times dose: $F_{3,42} = 2.841$, $p = .049$); there was also a main effect of dose ($F_{3,42} = 13.478$, $p < .001$). ACC-lesioned animals were not different from controls in any respect (group: $F_{1,14} = 1.661$, $p = .218$; lever \times group: $F < 1$, NS; dose \times group: $F_{3,42} = 2.043$, $p = .122$; lever \times dose \times group: $F_{3,42} = 1.2$, NS), even when the saline dose was considered on its own (lever: $F_{1,14} = 5.708$, $p = .032$; group: $F_{1,14} = 1.585$, NS; lever \times group: $F < 1$, NS).

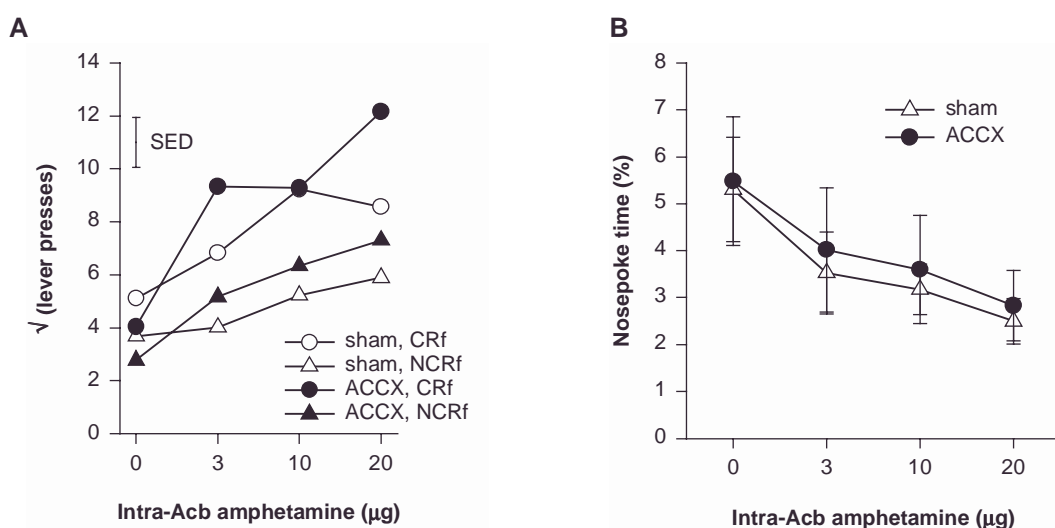


Figure 21. Responding for conditioned reinforcement, with intra-accumbens amphetamine. Lesions of the ACC had no effect on this task. **A**: Lever-pressing. *SED*, one standard error of the difference between means for the lever \times dose \times group term. **B**: Proportion of time spent nosepoking. Nosepokes during a CRf presentation were very few and were not included.

Nosepoking in the food alcove was dose-dependently reduced by intra-accumbens amphetamine, but this effect did not differ between groups (Figure 21B). An analysis by group \times (dose \times S) showed an effect of dose ($F_{2,563,35,886} = 9.571$, $\tilde{\epsilon} = .854$, $p < .001$), but no effect of group and no interaction ($F_s < 1$, NS).

Autoshaping

Data from one subject in the ACCX group (subject E3) were lost due to a malfunction, leaving 7 lesioned subjects and 10 sham-operated controls.

Acquisition

Lesioned animals were impaired at the acquisition of autoshaping (Figure 22). An analysis of difference scores revealed a significant impairment in the ACCX group (main effect of group, $F_{1,15} = 6.605$, $p = .021$), together with an effect of trial block ($F_{5,433,81,495} = 2.422$, $\tilde{\epsilon} = .604$, $p = .038$); the interaction was not significant ($F < 1$, NS). Analysis of ratio scores also demonstrated a significant impairment (group: $F_{1,15} = 8.966$, $p = .009$; trial block: $F_{5,066,75,984} = 1.475$, $\tilde{\epsilon} = .563$, NS; group \times trial block, $F < 1$, NS).

While sham subjects approached the CS+ faster than the CS-, lesioned rats approached the CS- faster than the CS+ (Figure 22D). Mean latencies to approach each stimulus were calculated across all trial blocks, and analysed using the model group \times (stimulus \times S), revealing a stimulus \times group interaction ($F_{1,15} = 7.295$, $p = .016$).

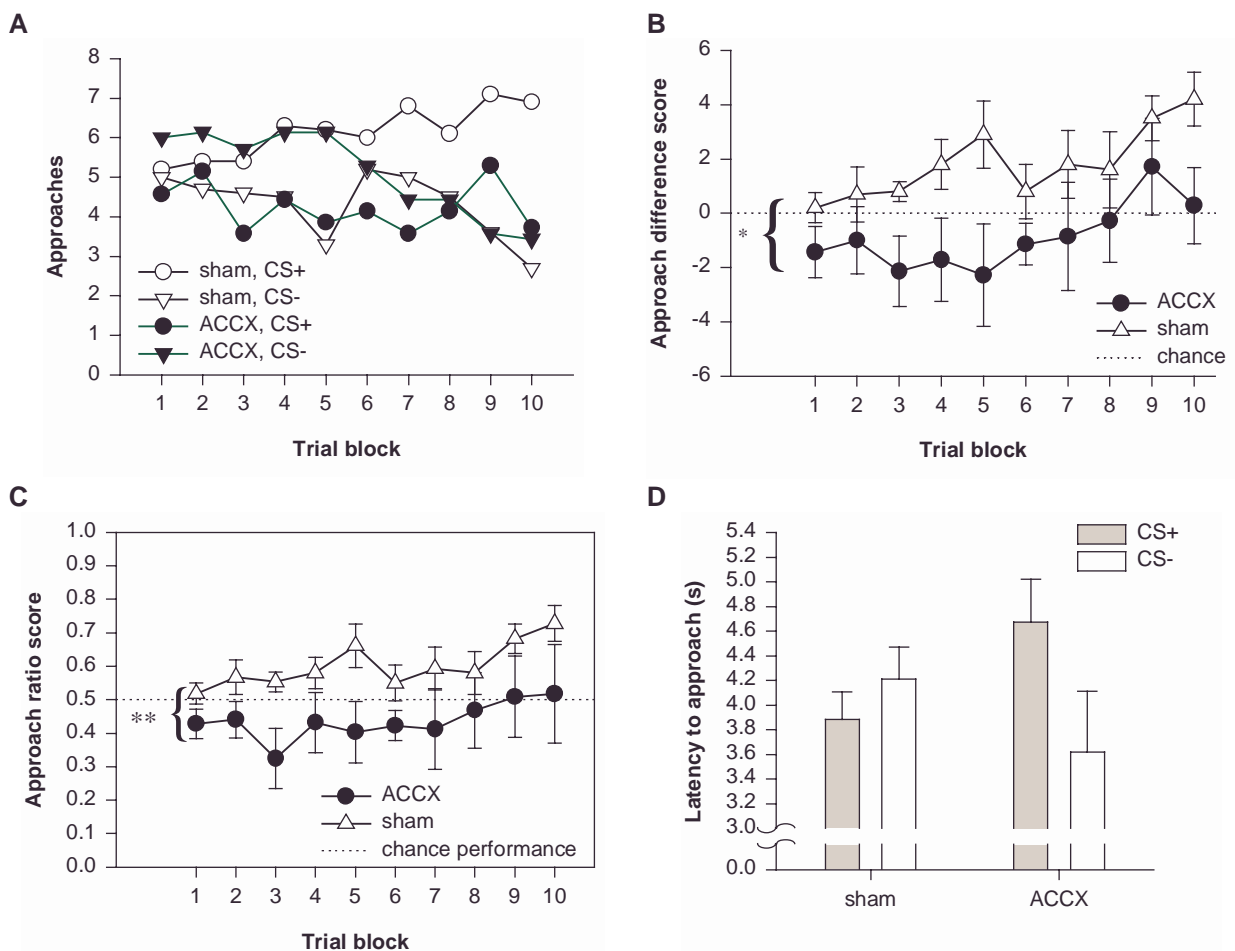


Figure 22. Autoshaping was impaired by lesions of the ACC. **A:** Approaches to the CS+ and CS- for each group. **B:** Approach data, expressed as a difference score (CS+ approaches - CS- approaches). **C:** Approach data, expressed as a discrimination ratio (CS+ approaches \div (CS+ approaches + CS- approaches)). **D:** Latencies to approach each stimulus, calculated across all trial blocks.

Probe test

In the probe test (Figure 23), there was a non-significant trend towards an impairment in the ACCX group. A discrimination ratio was calculated as the number of trials on which the CS+ was approached divided by the number of trials on which either stimulus was approached. This measure was analysed by one-way ANOVA, revealing no effect of group ($F_{1,15} = 3.928$, $p = .066$), even though the sham group discriminated between the stimuli (sham group compared to 50% discrimination ratio by one-sample t test: $t_9 = 5.673$, $p < .001$) and the ACCX group did not ($t_6 = 1.69$, $p = .142$).




Figure 23. Autoshaping probe test. Sham-operated controls approached the CS+ more than the CS- (as the number of approaches to the two stimuli are not independent, the proportion of trials on which the CS+ was approached was compared to 50%; *** $p < .001$). Though no such discrimination was detectable in the ACC-lesioned animals, the difference between groups did not reach significance ($p = .066$).

Omission training

Introduction of the omission contingency resulted in a reduction in the number of CS+ approaches, but the rate of reduction did not differ between groups (Figure 24). An ANOVA of the number of approaches to the CS+ for each trial block revealed a main effect of trial block ($F_{6,163,92.447} = 3.332$, $\tilde{\epsilon} = .685$, $p = .005$), and a main effect of group ($F_{1,15} = 5.06$, $p = .04$), reflecting the different starting points of the two groups, but no interaction ($F_{6,163,92.447} = 1.359$, NS).




Figure 24. Autoshaping omission test. The ACC-lesioned group approached the CS+ less than the sham group, and both groups' responding declined, but this difference was present from the start and the groups were not *differentially* affected by introduction of an omission contingency.

Sucrose consumption

Primary consummatory behaviour was unaffected by the lesion, with both groups consuming the same amount of sucrose (mean \pm SEM: ACCX 25.3 ± 2.1 ml, sham 27.7 ± 1.1 ml; $F_{1,16} = 1.056$, NS).

Locomotor activity in a novel environment

There was a trend towards hypoactivity in the ACC-lesioned group, but this failed to reach significance (Figure 25). An analysis of $\sqrt{(\text{beam breaks})}$ by group \times (bin \times S) revealed an effect of group that was close to significance ($F_{1,16} = 4.279$, $p = .055$), together with an effect of time bin ($F_{9,039,144.622} = 15.704$, $\tilde{\epsilon} = .822$, $p < .001$), reflecting habituation to the novel environment, with no interaction ($F < 1$, NS).




Figure 25. Locomotor response to novelty in sham- and ACC-lesioned rats.

Freezing to an aversive CS

Anterior cingulate-lesioned subjects did not differ from controls in their ability to freeze to a discrete CS predictive of footshock (Figure 26). An analysis of the percentage of time spent freezing, using the model group \times (stimulus presence \times S), showed no effect of group and no group \times stimulus interaction ($F_s < 1$, NS), despite a robust effect of the stimulus ($F_{1,12} = 429.856$, $p < .001$).




Figure 26. Freezing to an aversive CS+ was not affected by lesions of the ACC. The dependent variable is the percentage of time spent freezing, judged from video footage in 5-s bins. The 2 minutes preceding CS onset are compared with the 8 minutes following CS onset.

Summary

Lesions of the ACC did not affect subjects' ability to show temporally discriminated approach to a CS for food reward. This CS functioned successfully as a conditioned reinforcer in ACC-lesioned rats, and they showed normal potentiation of responding for conditioned reinforcement when given intra-accumbens amphetamine. They were not different from shams in measures of food consumption or locomotor activity, and were also capable of exhibiting conditioned freezing to an aversive CS. However, the same subjects were impaired at autoshaping.

Discussion

The present results establish that a substantial degree of Pavlovian conditioning can occur in rats with lesions of the ACC, although an autoshaping deficit was observed in the same animals, replicating previous findings (Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c). The implications will be discussed for each task used.

Temporally discriminated approach

ACC-lesioned animals were no different from sham-operated controls on any measure of temporally discriminated approach. This implies that, at the least, such animals can either form a Pavlovian association between the CS and the delivery of sucrose and use this representation to approach the CS, or can use the CS as a discriminative stimulus (S^D) for the performance of an instrumental approach response (the ambiguity as to whether this task measures Pavlovian or instrumental behaviour was discussed on p. 74). Inspection of Figure 20 (pp. 84/85) shows that the degree to which animals succeeded in approaching during the US directly paralleled the acquisition of responding to the CS. As the sucrose reward is only available for a brief time (5 s) in this task, it is obviously beneficial for the subjects to be nose-poking when the US begins; this illustrates the unavoidable S^D role of the CS.

Conditioned reinforcement

ACC-lesioned rats acquired an instrumental response with conditioned reinforcement, to the same level as controls. In this task, the response being tested has never had an instrumental relationship to food, so acquisition of discriminated lever-pressing demonstrates that the animals have acquired a Pavlovian association between the CS and some aspect of the food. In addition to leaving the efficacy of the conditioned reinforcer intact, the lesion did not impair the ability of intra-accumbens amphetamine to potentiate responding on the CRf lever, dose-dependently and selectively. Amphetamine also dose-dependently reduced the proportion of time the subjects spent nose-poking in the food/CS alcove (replicating a finding of Parkinson *et al.*, 1999b), perhaps because it potentiated the competing response of lever-pressing.

Strictly, of course, the present result is also explicable by a ‘novelty-seeking’ argument, also known as ‘sensory reinforcement’ (Kish, 1966) — the suggestion that animals work for the CS simply because it is interesting. However, this question has long since been addressed: Robbins & Koob (1978) demonstrated that a systemic dopamine indirect agonist, pipradrol, potentiated responding only for a CS explicitly paired with a primary reinforcer; this behavioural specificity has also been demonstrated for intra-accumbens amphetamine (Taylor & Robbins, 1984) and dopamine (Cador *et al.*, 1991).

As discussed earlier (p. 74), one suggested function of the ACC is to inhibit unrewarded responding. In the present study, ACC lesions did not increase approach during the unrewarded (VI) phase of the temporally discriminated approach task, or increase responding on the unrewarded (NCRf) lever in the conditioned reinforcement test. These data are therefore not compatible with the simple view that the ACC continuously suppresses responding that (on some occasions) leads to reward, although a role in inhibiting responding to unrewarded stimuli is not ruled out.

Autoshaping

The level of stimulus discrimination exhibited by ACC-lesioned animals in acquisition of the autoshaping task was significantly below that of control subjects, despite normal food consumption and locomotor behaviour in these animals. This result is especially noteworthy as the same animals were found to be unimpaired in the temporally discriminated approach task. At first glance, these tasks are extremely similar: both involve discriminated approach to a CS predictive of food reward. The two procedural vari-

ables that seem most likely to account for the difference are the location of the reward relative to the location of the CS (which are in the same location in the temporally discriminated approach task, and in separate locations in the autoshaping task) and the number of conditioned stimuli used (one versus two).

ACC-lesioned subjects also showed abnormal latencies to respond to the stimuli (as found by Bussey *et al.*, 1997a), and reduced discrimination in a probe test (though this difference was not significant). Though CS+/CS- discrimination was reduced in ACC-lesioned rats throughout training, this deficit was not precisely characterizable as an increase in CS- approaches, or a decrease in CS+ approaches; the former effect predominated early in training and the latter later on (Figure 22A, p. 86). Though clearly demonstrative of an impairment, the present study measured autoshaping in rats that already had considerable experience of CS-food pairings, and of lateralized responding (in the conditioned reinforcement test); for defining the autoshaping impairment more accurately, previous studies using naïve rats (Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c) may be more reliable.

The results of the omission test were not very informative. As the ACCX group approached the CS+ considerably less than the sham group at the end of omission training, and during the probe test, it was not surprising that they also did so at the beginning of the omission test. Both groups' responding declined during this test, but to the same degree. Although the groups were not differentially affected by the introduction of the omission contingency, the observation that their responding to the CS+ declined does not tell us a great deal. It would, of course, be expected that animals sensitive to the instrumental contingency would cease responding. However, a similar decline might be expected of a purely Pavlovian animal. Such an animal would initially respond to the CS, but by virtue of its responding, the US would not be presented and responding would extinguish (eventually to be followed by reinstatement of responding and extinction in a cyclical fashion; see Mackintosh, 1974, pp. 115/127).

Unconditioned measures of behaviour

Lesions of the ACC did not affect primary motivation or consummatory behaviour, as assessed by a sucrose consumption test. Similarly, the lesions did not significantly affect locomotor activity in a novel environment. This is one reason that the autoshaping deficit cannot be attributed to differences in general activity levels, the others being that a deficit was apparent even when considering CS+ approach as a proportion of those trials on which some stimulus was approached (the approach ratio score), and that absolute levels of responding in ACC-lesioned animals were comparable to those of sham-operated controls during the acquisition of autoshaping (Figure 22A, p. 86). There was a trend towards hypoactivity in the ACCX group, however, which is surprising given that Weissenborn *et al.* (1997) found a significant increase in the locomotor response to novelty in animals with ACC lesions. It may be that slight differences in lesion sites across the two experiments account for the difference (Weissenborn *et al.* used a post-geneal lesion; see Figure 14, p. 72).

Freezing to an aversive CS

ACC-lesioned rats exhibited normal conditioned freezing behaviour. The criterion used to judge freezing was strict, and it was apparent that following five CS→shock pairings, all animals were immobile for virtually the entire 8-min CS. In this experiment there were no unpaired controls, so it might be suggested that the freezing was an unconditioned response to the clicker CS; however, previous studies using exactly the same apparatus, stimuli, and assessment criterion as the present experiment have shown that freezing occurs at a level of ~20% when the clicker has been presented unpaired with shock, and ≥80% when paired (J. Hall, personal communication, March 1999; Hall, 1999).

These results may be contrasted to the demonstrations by Buchanan & Powell (1982a) and Gabriel *et al.* (Gabriel *et al.*, 1991a; Gabriel, 1993) that — in the rabbit — ACC lesions impair aversive Pavlovian conditioning and avoidance learning. Rather than appeal to procedural differences (the species difference, or the use of an aspirative lesion by Buchanan & Powell), the discrepancy may be explained through differences in the tasks used. Firstly, Buchanan & Powell observed normal eyeblink conditioning in their subjects, though heart-rate conditioning was impaired. As discussed in Chapter 1 (p. 40), aversive eyeblink conditioning is dependent upon the cerebellum; as Buchanan & Powell point out, even complete decortication does not prevent the acquisition of this CR (Oakley & Russell, 1972; 1975; 1976), and Gabriel *et al.* have shown a double dissociation between avoidance learning, which involves the ACC, and eyeblink conditioning, which does not (Steinmetz *et al.*, 1991; Gabriel *et al.*, 1996). It may be that freezing is another response that the ACC does not govern. Secondly, Buchanan & Powell found at least some heart-rate conditioning in ACC-lesioned rabbits, though the magnitude of cardiac deceleration was reduced compared to controls; Gabriel *et al.* have also reported acquisition of avoidance responding in rabbits with ACC lesions, though acquisition was retarded (Gabriel *et al.*, 1991a). Powell *et al.* (1994) found that although lesions of the ACC prevented rabbits from discriminating between a CS+ and a CS–, they did not abolish the conditioned response itself. Given the interesting dissociation in the present series of experiments between autoshaping and temporally discriminated approach tasks, discussed above, the necessity to discriminate between multiple stimuli may be a key factor in determining whether ACC lesions produce observable impairments in Pavlovian conditioning.

Summary

These data suggest that it is incorrect to characterize ACC-lesioned rats as being unable to form stimulus–reward associations. At some level, they are capable of Pavlovian conditioning, both appetitive and aversive. Nevertheless, lesions of the rat ACC have been clearly demonstrated to cause impairments in appetitive tasks that depend upon stimulus–reward associations, both in the autoshaping task used here and in previous studies (Bussey *et al.*, 1997a; Bussey *et al.*, 1997b; Parkinson *et al.*, 2000c). In what circumstances does this impairment in stimulus–reward learning manifest itself? This question will be addressed in Experiment 4.

EXPERIMENT 2: EFFECT OF ACC LESIONS ON AUTOSHAPING PERFORMANCE

Though ACC lesions have been repeatedly shown to disrupt the acquisition of autoshaping, it is presently unclear whether the ACC is involved in the storage of CS–US associations, in their behavioural expression, or in a learning process that regulates their formation. One way of addressing this issue is to examine the effects of lesions made *after* training. If the ACC is specifically involved in the formation of CS–US associations, a lesion should not affect performance in a well-trained animal. In contrast, such a lesion would be expected to disrupt performance if the ACC were critical for storage or retrieval of these associations. In the present study, rats were trained to an asymptotic level of performance on the autoshaping task before receiving lesions of the ACC and being re-tested.

Methods

Twenty-eight male hooded Lister rats were maintained at 90% of their free-feeding body mass and trained for 100 trials on the autoshaping task described earlier (p. 77). Subjects that failed to approach the CS+ on at least 70% of the last 30 trials were given a further 50 remedial trials; if they failed to meet the same criterion on the last 30 remedial trials, they were then excluded from the experiment. The successful subjects were given free access to food and randomly assigned to groups that received lesions of perigenual ACC or sham lesions; at the time of operation, they weighed 274–408 g. Following recovery, they were returned to the food deprivation regimen. Their performance on the same autoshaping task was tested for a further 50 trials; they then received a probe test (as described earlier) and 50 omission trials. After this a 2-h locomotor activity test was conducted in a novel environment with animals food-deprived.

Results

Eight subjects failed to reach the performance criterion (GL2, GL5, GL6, GL10, GL11, GL12, GL15, GL22). Of those that reached the criterion, 11 subjects received ACC lesions (GL1, GL3, GL4, GL9, GL17, GL18, GL19, GL20, GL21, GL23, GL24) and 9 subjects received sham lesions (GL7, GL8, GL13, GL14, GL16, GL25, GL26, GL33, GL34). There were two postoperative deaths in the sham group (GL7, GL14). Exploratory data analysis revealed that one subject in the sham group (GL16) was an outlier (complete absence of approach behaviour on two consecutive sessions with data points consistently >2 SD from the group mean); this subject was excluded from autoshaping and locomotor analysis. One subject in the lesioned group fell ill and was perfused after autoshaping was completed, but before locomotor testing. Histological analysis revealed that all lesions were correctly sited, so the final group sizes for the autoshaping performance test were 11 (ACCX) and 6 (sham).

Histology

In this group of ACC-lesioned subjects, neuronal loss and associated gliosis extended from ~ 2.5 mm anterior to bregma to ~ 0.3 mm posterior to bregma, destroying perigenual Cg1 and Cg2; as before, there was very slight damage to dorsal PrL in a few subjects and no damage to IL or PCC. Figure 27 presents schematics showing the largest and smallest extent of the lesions. (Photomicrographs of representative ACC lesions were shown in Figure 16, p. 80.)

Schematic of lesions




Figure 27. Lesions of the ACC (subjects GL1, GL3, GL4, GL9, GL17, GL18, GL19, GL20, GL21, GL23, GL24). Grey shading indicates the extent of the largest area of neuronal loss, and black the smallest. Diagrams are taken from Paxinos & Watson (1998).

Pre-operative acquisition

Both groups reached the same asymptote of performance pre-operatively (Figure 28). The difference scores for the last three blocks of training (10 trials per block) were analysed using the model group \times (block \times S); there were no significant effects of group ($F_{1,15} = 1.668$, NS), block ($F_{1.532,22.98} = 2.727$, $\tilde{\epsilon} = .766$, $p = .098$) or group \times block ($F < 1$, NS).

Post-operative performance

Lesions of the ACC significantly impaired the performance of the autoshaped response (Figure 28). Analysis of the post-operative difference scores using the model group \times (block \times S) revealed a significant main effect of group ($F_{1,15} = 7.765$, $p = .014$), reflecting poorer discrimination in the ACCX group; there was also a significant effect of block ($F_{4,60} = 3.524$, $p = .012$), but no interaction ($F < 1$, NS). This impairment was also evident following analysis of ratio scores, which also revealed an effect of group ($F_{1,15} = 5.73$, $p = .03$) and block ($F_{4,60} = 5.144$, $p = .001$). Although Figure 28(B,C) suggests some recovery in the ACCX group, there was no block \times group interaction ($F_{4,60} = 1.175$, NS).

Further analysis demonstrated that this deficit was attributable to a persistent deficit in CS+ approach in the ACCX group. Post-operative CS+ and CS- approach scores were analysed separately, using the model group \times (block \times S) in each case. These analyses showed that the ACCX group made significantly fewer approaches to the CS+ (main effect of group: $F_{1,15} = 5.221$, $p = .037$), an effect that did not alter across testing (terms involving block, $F_s < 1$, NS). The two groups did not differ in their approaches to the CS- (group: $F_{1,15} = 2.149$, NS); both groups showed an equivalent decline in CS- responding (block: $F_{4,60} = 6.462$, $p < .001$; block \times group: $F_{4,60} = 1.114$, NS). It is this decline in CS- responding that caused a degree of recovery of discriminative performance, evident as an improvement in difference and ratio scores, though the ACCX group remained impaired throughout testing.

Although these analyses did not demonstrate that the groups recovered at different rates, it was certainly the case that the ACCX group recovered to some extent (on all measures of performance), and did discriminate between the two stimuli. An improvement in discrimination scores was apparent for the ACCX group (main effects of block for difference scores, $F_{4,40} = 3.831$, $p = .01$; for ratio scores: $F_{4,40} = 5.922$, $p = .001$). Similarly, analysis of absolute approach scores in the ACCX group demonstrated a main effect of stimulus ($F_{1,10} = 28.495$, $p < .001$) and block ($F_{4,40} = 3.978$, $p = .008$) and a stimulus \times block interaction ($F_{4,40} = 3.831$, $p = .01$). No such improvement was detectable in the sham group, which performed well throughout (no block effects for the discrimination scores: maximum $F_{4,20} = 1.681$, NS; or stimulus \times block interaction for absolute approach scores: $F_{4,20} = 1.211$, NS).

Lesioned animals were slower to approach both stimuli. Mean latencies to approach each stimulus were calculated across all post-operative trial blocks, and analysed using the model group \times (stimulus \times S). This revealed a main effect of group ($F_{1,15} = 5.636$, $p = .031$), with the ACCX group showing longer approach latencies, and a main effect of stimulus ($F_{1,15} = 15.543$, $p = .001$) as subjects approached the CS+ faster than the CS-. There was no stimulus \times group interaction ($F_{1,15} = 1.413$, NS).

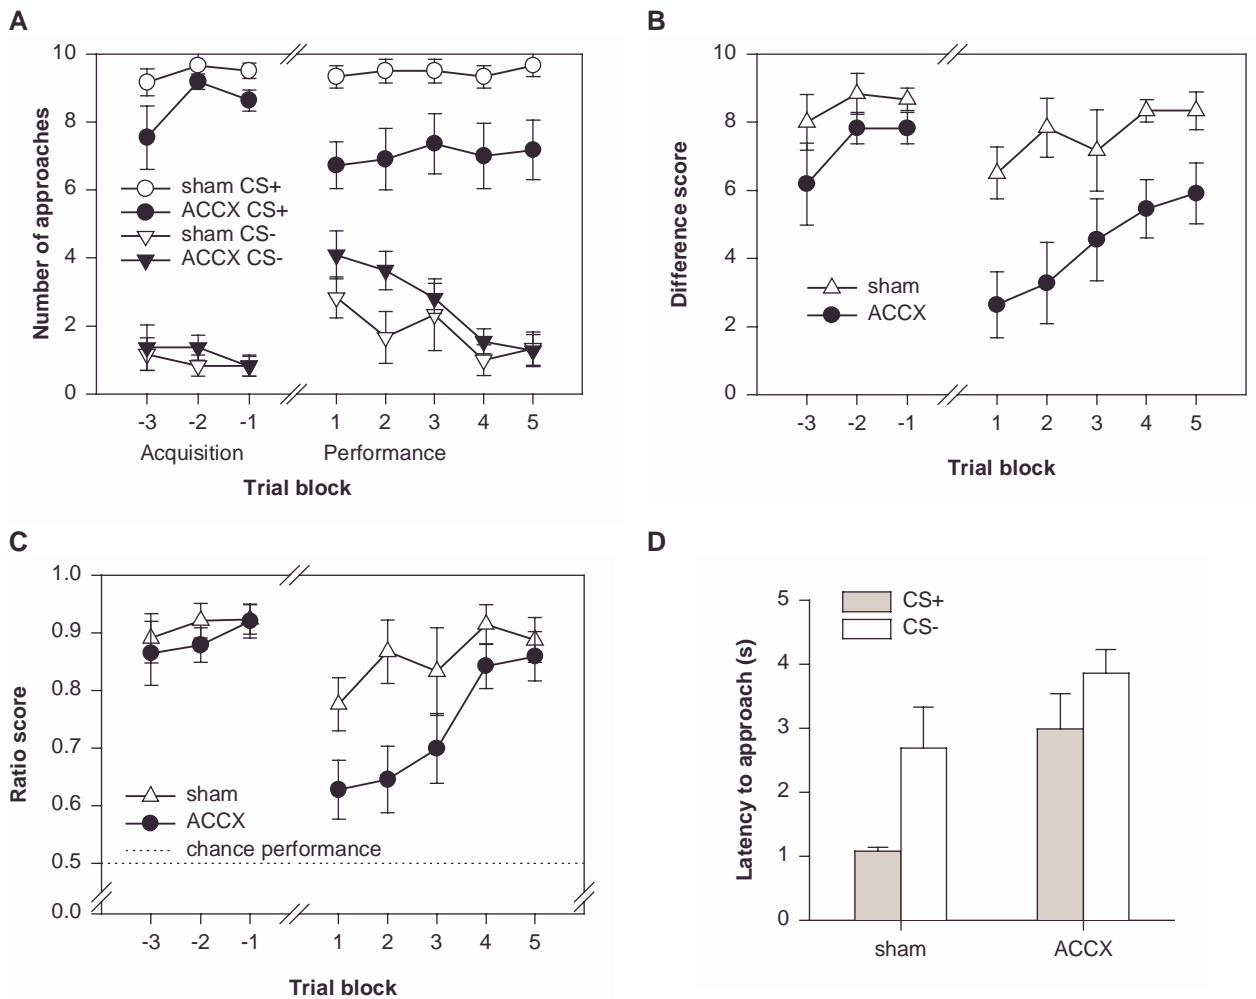


Figure 28. Lesions of the ACC impair the performance of autoshaping when those lesions are made after training. The final three sessions of pre-operative acquisition training are shown, together with post-operative performance. **A:** Approaches to CS+ and CS-. **B:** Difference scores. **C:** Discrimination ratio scores, calculated as for Figure 22 (p. 86). **D:** Latencies to approach each stimulus post-operatively. The ACCX group approached more slowly.

Probe test

ACC-lesioned subjects showed reduced discrimination in the probe test (Figure 29). A discrimination ratio was calculated as the number of trials on which the CS+ was approached divided by the number of trials on which either stimulus was approached. Analysis of this measure by one-way ANOVA revealed an impairment in the ACCX group ($F_{1,15} = 4.566$, $p = .049$). However, both groups discriminated between the CS+ and CS- (sham group compared to 50% discrimination ratio by one-sample t test: $t_5 = 22.077$, $p < .001$; ACCX group: $t_{10} = 9.515$, $p < .001$).




Figure 29. Autosshaping probe test, showing reduced CS+/CS- discrimination in ACC-lesioned rats. (N.B. Approaches to the CS+ and CS- are mutually exclusive on any given trial.)

Omission test

Lesioned subjects did not differ from shams in their response to the introduction of an omission contingency (Figure 30). Analysis of the number of trials on which the CS+ was approached using the model $\text{group} \times (\text{block} \times \text{S})$ revealed a near-significant effect of group ($F_{1,15} = 4.269$, $p = .057$), reflecting the previously-established lower baseline of CS+ approaches in the ACCX group, but no effect of block ($F_{4,60} = 1.226$, NS) and, critically, no block \times group interaction ($F < 1$, NS). (It may be worth noting that even if the ACCX group had ceased responding to the CS+ more rapidly, for which there was no statistical proof, it could not be stated with confidence that they were ‘more instrumental’ animals, better able to inhibit Pavlovian conditioned responding, as more rapid Pavlovian extinction would be an alternative explanation.)




Figure 30. Autoshaping omission test; ACC-lesioned rats did not differ significantly from sham-operated controls.

Locomotor activity in a novel environment

As one animal in the ACCX group (rat GL3) fell ill after the autoshaping tests and was perfused, final group numbers for the locomotor test were 10 (ACCX) and 6 (sham).

There was no clear pattern of difference in locomotor activity between sham and ACCX groups, although there were statistical differences in the pattern of habituation to novelty. Beam-break data were subjected to a square-root transformation and analysed using the model $\text{group} \times (\text{bin}_{12} \times \text{S})$. There was no main effect of group ($F < 1$, NS), but in addition to the main effect of bin ($F_{8,49,118.859} = 27.459$, $\tilde{\epsilon} = .772$, $p < .001$), reflecting habituation, there was a bin \times group interaction ($F_{8,49,118.859} = 3.48$, $\tilde{\epsilon} = .772$, $p = .001$). The only bin for which a simple effect was significant in its own right was the bin finishing 70 min into the session (simple effect of group for this bin, $F_{1,14} = 5.205$, $p = .039$), but elimination of this bin left the interaction term still significant ($F_{7,628,106.793} = 2.883$, $\tilde{\epsilon} = .763$, $p = .007$). Inspection of Figure 31 suggests that this may have been due to slight hyperactivity in the ACCX group late in the session.

Summary

Lesions of the ACC impaired the performance of an autoshaped response that had been trained to asymptote pre-operatively. This deficit was primarily due to a decrease in CS+ approaches in the lesioned animals, and persisted throughout testing.




Figure 31. Locomotor response to novelty. 120-min session scored in 10-min bins.

Discussion

Nature of the autoshaping deficit

Previous studies have consistently found ACC-lesioned rats to approach the CS– more than control rats (Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c), whereas the performance deficit in the present experiment was due to a reduction in CS+ approach (the transient post-operative increase in CS– approach was not significantly greater than that observed in sham-operated controls). It may be that a failure to discriminate between the two stimuli may naturally manifest itself as an alteration in either decreased CS+ or increased CS– responding, influenced by the level of general activity of the subjects in the autoshaping apparatus, a factor to which the slight differences in lesion coordinates may have contributed. The most anterior injection in the present experiments was advanced rostrally by 0.4 mm compared to previous autoshaping studies (Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c), and by 0.5 mm compared to Weissenborn *et al.* (1997), the only comparable study for which locomotor data are available. Indeed, while Weissenborn *et al.* (1997) found their ACC-lesioned rats to be clearly hyperactive, this was not obvious in the present experiments (Figure 25, p. 88; Figure 31, p. 97). In further support of this interpretation, reduced CS+ approach, as well as reduced CS– approach, was observed during acquisition in Experiment 1 (Figure 22A). Nevertheless, impairments in CS+/CS– discrimination have been a consistent feature of autoshaping in ACC-lesioned rats, whether assessed by acquisition or performance testing, or by probe tests (Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c; present data).

In the present study, a very high level of CS+ approach was attained pre-operatively, compared to previous studies (Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c); this was entirely due to the use of a selection criterion. Another interpretation supported by the present data, therefore, is that the ACC contributes to CS+ approach especially in those animals who discriminate very well; this effect may usually be masked by variability in normal subjects.

Speed of responding

Lesioned subjects were slower to approach the stimuli. This accords with the observations of Bussey *et al.* (1997a), but not Parkinson *et al.* (2000c), who found that ACC-lesioned rats approached the stimuli faster than control subjects. Again, it is possible that slight differences in lesion location account for these discrepancies. Indeed, Parkinson *et al.* (2000c) suggested that ventral ACC damage was more likely to produce stimulus–reward learning deficits while dorsal ACC damage might affect the speed of respond-

ing, based on the distinction between the dorsal, 'motor' ACC and the ventral, 'emotional' ACC (Dum & Strick, 1993; Neafsey *et al.*, 1993). The present lesions were slightly more anterior than those of Bussey *et al.* (1997b), and damaged dorsal ACC more consistently than those of Parkinson *et al.* (2000c), and, as suggested above, there is some indirect evidence (via Weissenborn *et al.*, 1997) that the ACC-lesioned subjects in the present study were not as active as those in previous autoshaping studies that used more posterior coordinates, though this cannot be established directly. However, if more anterior, more dorsal ACC lesions produce less hyperactivity, it is not clear why Parkinson *et al.* (2000c) found more rapid responding in ACC-lesioned rats, given that their lesions were slightly more anterior on average than those of Bussey *et al.* (1997a).

Inspection of the data from these studies suggests that the approach latencies of ACC-lesioned rats are actually more consistent across studies than those of the sham-operated controls, in terms of absolute magnitude (all these studies used the same apparatus). Another hypothesis, therefore, is that the selection process contributed to the observed results: the selection of subjects with especially high levels of discrimination led to unusually low approach latencies in the sham group. This hypothesis would suggest that the ACC contributes to rapid approach in animals who discriminate well. However, it is not clear how this hypothesis would account for the significant *reduction* of approach latency in ACC-lesioned rats found by Parkinson *et al.* (2000c). It has been suggested that autoshaping exhibits important individual variability (Tomie *et al.*, 1998; 2000); it would be intriguing if the ACC is a source of this variability. If applied to approach latencies, this hypothesis would predict that the variance of ACC-lesioned, unselected rats was smaller than the variance of sham-lesioned, unselected rats. There have been no direct tests of this hypothesis, as variance comparisons have not been published for the present autoshaping task; however, inspection of Figure 22D (based on non-naïve subjects; page 86) does not support this suggestion.

Learning versus performance

The finding that post-training ACC lesions impair the performance of autoshaping strongly suggests that its role is not limited to learning the stimulus–reward associations; instead, the ACC is involved in storage or retrieval of the associations (a mnemonic role), or in the mechanism of behavioural expression.

This is a difficult question to answer conclusively, as it may be argued that further pre-operative training might have rendered performance independent of the ACC, or that recovery might eventually have been observed post-operatively in ACC-lesioned subjects. There is some evidence that the ACC has a time-limited role (that is, the ACC is particularly important *early* in learning). Gabriel *et al.* (1991a) have reported that two-way active avoidance behaviour is eventually acquired in ACC-lesioned rabbits. Furthermore, when ACC lesions, or combined lesions of MD and anteroventral (AV) thalamus, are made after the acquisition of avoidance behaviour, the lesions do not impair performance as much as they do when made before training (Gabriel *et al.*, 1980a, p. 162; Gabriel, 1993; Freeman *et al.*, 1996; Hart *et al.*, 1997). They suggest that the ACC (with the MD thalamus) rapidly acquires the CS+/CS– discrimination, then 'teaches' the AV and the PCC, 'relegating' the discriminative role and releasing the ACC for further learning; Gabriel *et al.* have even suggested this as an analogue of behavioural automatization (Gabriel *et al.*, 1980a, pp. 143–163 / 220; Freeman *et al.*, 1996). In the words of Hart *et al.* (1997), the engram is nomadic. In accordance with this hypothesis, lesions of AV or PCC did not impair acquisition, but impaired asymptotic performance of this task (Gabriel *et al.*, 1983; Gabriel *et al.*, 1987). To some degree, this would also accord with views of primate ACC as a specialized 'error detector' (to be discussed later); a structure responsible for error detection and correction might become less important as the task is automatized. It is not known whether such an 'automatization' account is applicable to autoshaping, a putative Pavlovian response; this suggestion could be envisaged as the supersession of CS–US associations by

CS–UR associations as the dominant representation controlling behaviour, a hypothesis that would be testable by devaluing the US after brief or extended training.

There is some evidence from rat studies to support a role for the ACC early in learning. Parkinson *et al.* (2000c) found that a degree of discrimination was eventually acquired by ACC-lesioned subjects in the autoshaping task, though they never attained the performance of shams (see Parkinson *et al.*, 2000c, p. 49). Bussey *et al.* (1996) found that ACC lesions facilitated early learning of a CVD, a task that may depend on S–R associations, while PCC lesions impaired late learning. These results and those of Gabriel's group (see Freeman *et al.*, 1996) would be anticipated if the ACC formed stimulus–reward associations early in training, while the PCC formed stimulus–response habits (Bussey *et al.*, 1996; 1997a; 1997b). In tasks where both systems contribute, such as autoshaping and Gabriel's active avoidance task, ACC lesions impair performance early in acquisition (Gabriel, 1990; Gabriel *et al.*, 1991a; Gabriel, 1993; Parkinson *et al.*, 2000c), while PCC lesions impair late performance of active avoidance (Gabriel *et al.*, 1987). In S–R tasks, where the two systems compete, ACC lesions improve performance early in acquisition (Bussey *et al.*, 1996) while PCC lesions impair performance later on (Bussey *et al.*, 1996; 1997b). In tasks only soluble via stimulus–reward associations, PCC lesions may improve performance (Bussey *et al.*, 1997b). A critical prediction of the hypothesis of Gabriel *et al.* (that the ACC teaches the PCC) is that PCC lesions, which do not impair the acquisition of autoshaping (Bussey *et al.*, 1997a), would impair performance if the response were overtrained.

In support of this compelling account, there was partial recovery of the ACCX group in the present experiment, and they did discriminate between the stimuli, both in the performance test and the probe test. However, the ACCX group did not recover fully. The recovery was largely due to a decline in CS–responding; the deficit in CS+ approach was persistent and showed no signs of recovery. As these animals were not hypoactive in a locomotor test, there is no reason to think that the deficit in CS+ responding was due to a general lack of activity. Though it may be that the autoshaping response was not sufficiently overtrained to observe normal function after ACC lesions (compare Hart *et al.*, 1997), the response was behaviourally asymptotic before the lesion.

EXPERIMENT 3: EFFECTS OF ACC LESIONS ON ‘SIMPLE’ PAVLOVIAN-INSTRUMENTAL TRANSFER

Pavlovian conditioned stimuli may elicit autonomic or skeletomotor conditioned responses, and serve as behavioural goals (conditioned reinforcers), but may also elicit conditioned ‘motivational’ responses, as discussed in Chapter 1 (p. 26). One example is Pavlovian–instrumental transfer (PIT), in which an appetitive Pavlovian CS potentiates ongoing instrumental responding. The results of Experiment 1 demonstrated that the ACC is not necessary for simple Pavlovian conditioning; accordingly, it was anticipated that normal PIT should be observed in ACC-lesioned rats.

In this task, a Pavlovian association is first established between a CS and reward. Subjects are then trained to respond instrumentally for the same reward (with no CS present), and in an extinction test, responding is assessed in the presence and absence of the CS. This is a ‘simple’ test of PIT, in that the instrumental reinforcer is the same as the Pavlovian US (see Chapter 1).

Methods

Subjects

The subjects were those that previously served in the autoshaping performance study, except for two animals that fell ill after locomotor testing and were perfused (rats GL4 and GL8). This left $n = 6$ (sham) and 9 (ACCX). These subjects were tested on a Pavlovian–instrumental transfer task of the simple kind, described below.

Simple Pavlovian to instrumental transfer

The task was conducted in the standard operant chambers, which were new to the subjects. The method was based on Balleine (1994).

Throughout the experiment, the reinforcer used was one 45-mg sucrose pellet (Rodent Diet Formula P, Noyes, Lancaster, NH). The task used two stimuli. Stimulus 1 consisted of the left and right stimulus lights (2.8 W bulbs) flashed at 3 Hz. Stimulus 2 was a clicker relay operated at 10 Hz. These stimuli were designated the CS and neutral stimulus (NEUT) in counterbalanced fashion. A 2.8-W houselight was illuminated throughout.

Pavlovian training. Eight training sessions were given. Each session contained six 2-min presentations of the CS, during which reinforcement was delivered on a random time (RT) 30-s schedule. Stimulus presentations were separated by an interstimulus interval (ISI) of 2–4 min, during which no reinforcement was given. Conditioning was assessed as a discrimination ratio: the proportion of total nose-poking time during the CS, corrected for the differences in CS and ISI duration (that is, $CS\% / \{CS\% + ISI\% \}$). In the final session, two 2-min presentations of the NEUT stimulus were also given, unreinforced, to reduce unconditioned suppression when this stimulus was subsequently presented during the test phase.

Instrumental training. Instrumental training was conducted in eight 30-min sessions with a single lever present. Responding was reinforced on a random interval (RI) schedule, whose parameter in subsequent sessions was 2, 15, 30, and thereafter 60 s.

Instrumental extinction. A single 30-min session was given in which the lever was available but unreinforced, following the observation that PIT is best observed when the response has been partially extinguished (Dickinson *et al.*, 2000, p. 473). No further Pavlovian sessions were given after instrumental training.

Transfer test. The transfer test was conducted over two sessions with the lever present but never reinforced. In each session, the CS, NEUT, and ISI were presented four times each; the stimuli (including the ISI) all lasted 2 min and were randomized in triplets, with the constraint that the same stimulus was never presented in two consecutive 2-min periods.

Results

Pavlovian training

The sham and ACCX groups did not differ in their stimulus-related behaviour during Pavlovian training (Figure 32). The approach ratio during Pavlovian sessions was calculated from the proportion of the CS spent nose-poking (%CS) and the proportion of the ISI spent nose-poking (%ISI) as follows: approach ratio = (%CS) ÷ (%CS + %ISI). As pellets were being delivered during CS presentation, this measure is not a pure measure of conditioned responding, being contaminated by unconditioned approach to the food. However, the two groups did not differ: an analysis using the model group × counterbalancing × (session × S) revealed no effect of group ($F_{1,11} = 2.023$, NS) and no group × session interaction ($F_{7,77} = 1.045$, NS), with the main effect of session approaching significance ($F_{7,77} = 1.981$, $p = .068$). Subjects nose-poked more during the clicker than the light CS (mean approach ratios 0.681 and 0.603, respectively; main effect of counterbalancing: $F_{1,11} = 6.555$, $p = .027$) but there were no other effects of the counterbalancing condition ($F_s < 1$, NS).




Figure 32. Pavlovian training. The approach ratio is the proportion of total nose-poking behaviour occurring at times when the CS was presented (see text). Both groups approached the alcove more during the CS than during the interstimulus interval, with no group differences. As food was delivered during the CS, the approach behaviour partly reflects unconditioned responding.

Instrumental training

Both groups acquired the instrumental response at the same rate (Figure 33). Lever-press data from instrumental acquisition sessions were subjected to a square-root transformation and analysed using the model group × (session × S). There was no effect of group, and no group × session interaction ($F_s < 1$, NS), though there was a main effect of session ($F_{4,316,56,114} = 11.528$, $\tilde{\epsilon} = .617$, $p < .001$). Similarly, responding did not differ between the groups during the extinction session (univariate ANOVA, $F < 1$, NS).




Figure 33. ACC lesions did not impair the acquisition of a free-operant instrumental response, or affect responding in extinction (*ext*, extinction session).

Transfer test

The CS reliably elevated responding relative to the ISI and the neutral stimulus, and this effect did not differ between groups (Figure 34). Response rates for the two test sessions were square-root transformed and analysed using the model $\text{group}_2 \times \text{counterbalancing}_2 \times (\text{session}_2 \times \text{stimulus}_3 \times S)$, where stimulus had three levels (CS, ISI and NEUT) and counterbalancing had two (light or clicker CS). Predictably, subjects responded more on the first test session than the second (effect of session: $F_{1,11} = 74.968$, $p < .001$), but there were no other effects of the test session. Similarly, the counterbalancing condition had no effect on responding; thus, the light and clicker were equally effective as CSs. The CS significantly affected behaviour (stimulus: $F_{2,22} = 72.784$, $p < .001$). Pairwise comparisons using a Sidak correction showed that responding during the CS was greater than during the ISI or the NEUT stimulus ($p < .001$), which did not differ from each other ($p = .966$). The sham and ACCX groups did not differ in any respect (maximum $F_{2,22} = 1.549$, NS).




Figure 34. Transfer test. ACC lesions did not affect Pavlovian–instrumental transfer; the CS elevated responding relative to the ISI and a neutral stimulus.

Discussion

These results provide a further demonstration of normal Pavlovian conditioning in ACC-lesioned rats, who exhibited normal PIT, indicating that the conditioned motivational impact of the appetitive CS (see Dickinson, 1994) was intact and able to modulate instrumental behaviour. In addition, it demonstrates normal free-operant instrumental acquisition in ACC-lesioned rats. Finally, there was no evidence that the ACCX group generalized to the neutral stimulus, which was of a different modality to the CS. If the ACC is indeed critical for discriminating between rewarded and unrewarded stimuli, as discussed earlier, it must be assumed that the visual and auditory CSs were too different for generalization to have occurred in the first place. However, the idea that the rat ACC is necessary to discriminate similar stimuli on the basis of their association with reward (but is not required for conditioning *per se*) has not yet been tested directly. Experiment 4 will do so.

EXPERIMENT 4: EFFECTS OF ACC LESIONS ON A TWO-STIMULUS DISCRIMINATED APPROACH TASK

Experiment 1 demonstrated a striking dissociation in which ACC-lesioned rats successfully learned to approach a single appetitive CS in a temporally discriminated approach task but were impaired at auto-shaping. Indeed, the neural basis of these tasks has been dissociated before: CeA lesions impair auto-shaping (Parkinson *et al.*, 2000b) but not temporally discriminated approach (Robledo *et al.*, 1996). Therefore, a further experiment was designed to explore the difference between the two tasks. As discussed earlier, these two tasks differ in two main ways.

The first is the location of the CS relative to the US. In the temporally discriminated approach task, the CS is presented in the same spatial location as the food, while in the autoshaping task, approach to the CS takes the subject away from the food source. It may be that the ACC is critical for appetitive approach to a CS, but not for approach to a US (literally, sign-tracking versus goal-tracking, or preparatory versus consummatory behaviour).

This might also reflect the differential contribution of Pavlovian and instrumental responding. As discussed in Chapter 1 (p. 41), autoshaping is most probably a Pavlovian response — the alternative explanation, that it reflects instrumental approach to a conditioned reinforcer (Williams, 1994a), cannot easily explain the ACC impairment, as it has now been shown that ACC-lesioned animals can work normally for a CRf. In the temporally discriminated approach task, there is an unavoidable instrumental contingency between approach to the site of the CS and food acquisition: the CS might serve as a discriminative stimulus for instrumental approach. However, a version of the discriminated approach task with an inadvertent instrumental contingency has also been employed; in this version (Burns *et al.*, 1993), nosepokes caused the sucrose dipper to rise (B.J. Everitt, personal communication, 29 January 1999). Different effects of BLA lesions have been observed on the Pavlovian and instrumental versions of this task (Cador *et al.*, 1989; Burns *et al.*, 1993), which weakens the argument that the Pavlovian version used in Experiment 1 differs *further* from autoshaping in the degree of instrumental contingency. Nevertheless, this remains a possibility.

In summary, this difference between the two tasks leads to the hypothesis (Hypothesis 1) that the rat ACC is critical for Pavlovian conditioned approach, not instrumental or consummatory approach behaviour, and not other simple forms of Pavlovian conditioning (such as conditioned freezing or PIT).

The second difference is the number of stimuli used. In the autoshaping task, the subject is required to discriminate two stimuli, identical apart from their location. In the simple discriminated approach task, the discrimination is temporal: the subject is merely required to discriminate the presence of a single stimulus from its absence. The hypothesis that follows from this (Hypothesis 2) is that the rat ACC is necessary for discriminating similar stimuli on the basis of their association with reward.

To distinguish these two possibilities, a task was designed that had features of both the temporally discriminated approach and autoshaping tasks. Approach was to the food source, as in the temporally discriminated approach task, but two similar stimuli governed approach, as in autoshaping. One stimulus (CS+) signalled the imminent delivery of sucrose solution to a food alcove, while the other (CS-) did not. Essentially, this task is identical to autoshaping except that approach is *measured* to the food alcove, rather than to the stimuli. Pilot experiments established that normal rats could discriminate the two stimuli, although with difficulty. Finding an impairment in ACC-lesioned rats with this task would therefore support Hypothesis 2, and normal performance would support Hypothesis 1. In addition, a conditioned reinforcement test was given using the two stimuli.

Methods

Overview

Naïve subjects received lesions of the ACC ($n = 12$) or sham lesions ($n = 12$); their body mass at the time of surgery was 333–379 g. Following recovery, they were maintained at 85% of their free-feeding mass. The subjects were subsequently trained for 12 sessions on a two-stimulus discriminated approach task (described below), as pilot studies had determined that significant CS+/CS– discrimination emerged in normal animals within this time; a conditioned reinforcement test was then conducted for two sessions.

Two-stimulus temporally discriminated approach task

This task was conducted in the operant chambers used for the temporal discriminated approach task described previously (p. 76).

The levers were not extended during training. The stimulus lights located above the levers were designated the CS+ and CS–, counterbalanced left/right across rats. At the start of every session, the houselight was on and the dipper was lowered. This phase lasted for a VI of 30–90 s. Next, the houselight was extinguished and one of the stimulus lights was illuminated for 5 s. Following presentation of the CS+, the houselight was illuminated and the dipper raised for 5 s to deliver 10% sucrose solution; this constituted the US. Following presentation of the CS–, the houselight was similarly illuminated but the dipper was not raised, and a brief click was generated in order that both stimuli had an auditory and a visual component. Regardless of the stimulus, the chamber was then in the starting state and the next VI began.

One trial consisted of a presentation of the CS+ and a presentation of the CS–; the order of the stimuli was randomized within each trial. A session consisted of 15 trials, after which the houselight was extinguished. Subjects received one session per day. For each period (VI, CS+/CS–, US or a notional 5-s equivalent following the CS–), the number of alcove entries and the time spent nose-poking in the alcove were recorded.

Two-stimulus test of conditioned reinforcement

This task was conducted in the same operant chambers. Two 30-min sessions were given on consecutive days, during which the houselight was illuminated and two levers were available, designated the CRf and NCRf levers. Responding on the CRf lever produced an abbreviated version of the CS+ with probability 0.5, while responding on the NCRf lever produced an abbreviated version of the CS– with probability 0.5. The abbreviated CS+ was produced by extinguishing the houselight and illuminating the CS+ stimulus light for 0.5 s, after which the houselight was re-illuminated, the stimulus light was switched off and the empty dipper was raised for 0.3 s. The corresponding CS– stimulus was identical except that the other stimulus light was used, and a click replaced elevation of the dipper. The levers were assigned so that the CRf lever was located underneath the CS+ stimulus light, and the NCRf lever under the CS– stimulus. Lever-pressing and nose-poking were recorded in 5-min bins.

Results

Histology

Histological analysis determined that two of the lesions in the ACCX group were incomplete, and these subjects (I2, I11) were excluded. Neuronal loss and associated gliosis extended from ~2.7 mm anterior to bregma to ~0.3 mm posterior to bregma. However, the ACCX group was somewhat heterogeneous; 4 animals had lesions including the ventral perigenual portion of Cg2 at 1.6–1.7 mm anterior to bregma (subjects I5, I6, I9, I12; Figure 36 presents a schematic showing the largest and smallest extent of the lesions in these subjects), while 6 animals had lesions that did not extend this far ventrally (subjects I1, I3, I4, I7, I8, I10; schematics of lesions in these subjects are shown in Figure 35). As retrograde tracing studies (Parkinson, 1998) have indicated that this region of the ACC projects most strongly to the AcbC,

strongly implicated in appetitive approach behaviour (Parkinson, 1998; Parkinson *et al.*, 1999c; Parkinson *et al.*, 2000c), analyses were conducted using both the complete lesion group (ACCX group, $n = 10$) and the subgroup with ventral perigenual lesions (designated the ACCX-whole group, $n = 4$). No sham animal was excluded, leaving $n = 12$ for this group (subjects I13, I14, I15, I16, I17, I18, I19, I20, I21, I22, I23, I24). Photomicrographs of representative ACC lesions were shown in Figure 16, p. 80.

Schematic of lesions




Figure 35. Lesions of the ACC excluding the ventral perigenual region. Grey shading indicates the extent of the largest area of neuronal loss, and black the smallest. Diagrams are taken from Paxinos & Watson (1998). Subjects: I1, I3, I4, I7, I8, I10. Subjects were classified as having whole or partial ACC lesions on the basis of whether the ventral portion of Cg2 in the 'cup' of the genu was lesioned (seen here in sections +1.6 and +1.7 mm from bregma).




Figure 36. Lesions of the ACC including the ventral perigenual region. Grey shading indicates the extent of the largest area of neuronal loss, and black the smallest. Diagrams are taken from Paxinos & Watson (1998). Subjects: I5, I6, I9, I12. Compare Figure 35.

Two-stimulus discriminated approach task

As this task was designed to be comparable to the autoshaping task used previously, but also to the temporally discriminated approach task, two primary measures of performance were used.

Firstly, for direct comparison with autoshaping, the number of trials was calculated on which at least one nosepoke occurred during stimulus presentation, for both the CS+ and the CS-. From these, difference and ratio scores were calculated, as for the autoshaping task. (If no approach occurred to either stimulus during a session, a ratio score of 0.5 was assigned, though this was a very rare occurrence.)

Secondly, for comparison with previous temporally discriminated approach tasks, an approach discrimination ratio was calculated: the *proportion* of each stimulus period spent nosepoking (%stimulus) was compared to the proportion of the ISI spent nosepoking (%ISI) using the formula *discrimination ratio* = %stimulus ÷ (%stimulus + %ISI). This ratio was calculated for both the CS+ and CS-, and ISI responding was calculated over both ISI periods in the corresponding trial (including the ISI preceding the CS+ and that preceding the CS-). Therefore, the ratios for CS+ and CS- are directly comparable, as both are calculated relative to the same %ISI.

Analyses based on the number of trials on which approach occurred

The ACCX group were impaired in their ability to discriminate between the two stimuli (Figure 37).

Analysis of absolute approach scores using the model $group_2 \times (stimulus_2 \times session_{12} \times S)$ demonstrated that the ACCX group made fewer approaches overall (main effect of group: $F_{1,20} = 7.48, p = .013$). There was a main effect of stimulus ($F_{1,20} = 57.626, p < .001$), and of session ($F_{5,516,110.326} = 53.52, \tilde{\epsilon} = .501, p < .001$), and a stimulus \times session interaction ($F_{11,220} = 14.443, p < .001$). In addition, there were stimulus \times group ($F_{1,20} = 6.827, p = .017$) and stimulus \times session \times group ($F_{11,220} = 2.178, p = .017$) interactions. The session \times group interaction was not significant ($F < 1, NS$).

This highly complex pattern of results was investigated using simple effects analyses. First, the CS+ and CS- were considered separately. The ACCX group responded less to the CS+ than the sham group (group: $F_{1,20} = 9.567, p = .006$) across all sessions (session: $F_{7,952,159.048} = 60.114, \tilde{\epsilon} = .723, p < .001$; session \times group, $F = 1.024, NS$). The ACCX group also responded less to the CS- than did shams (group: $F_{1,20} = 4.458, p = .048$), again in a session-independent manner (session: $F_{6,658,133.153} = 24.454, \tilde{\epsilon} = .605, p < .001$; session \times group: $F < 1, NS$). Second, the ACCX and sham groups were considered separately. The sham group learned to discriminate between the stimuli (stimulus: $F_{1,11} = 56.89, p < .001$; session: $F_{4,978,54.759} = 31.375, \tilde{\epsilon} = .453, p < .001$; stimulus \times session: $F_{10,555,111.106} = 7.076, \tilde{\epsilon} = .96, p < .001$). The ACCX group also learned to discriminate (stimulus: $F_{1,9} = 11.452, p = .008$; session: $F_{5,137,46.233} = 23.142, \tilde{\epsilon} = .467, p < .001$; stimulus \times session: $F_{11,99} = 11.085, p < .001$). Third, the groups' performance was considered for each session. The ACCX group showed discrimination between CS+ and CS- ($p < .05$) from session 9 on, while the sham group first showed discrimination on session 4 (and subsequently on sessions 6 and 8–12).

These analyses indicate that both groups acquired discrimination, with the shams acquiring faster, but do not answer the question of whether the *degree* of discrimination differed between groups. For this, direct measures of discriminative ability were used.

Analysis of difference scores (approaches during the CS+ – approaches during the CS-) using the model $group_2 \times (session_{12} \times S)$ revealed a significant main effect of group ($F_{1,20} = 6.827, p = .017$). In addition, there was a main effect of session ($F_{11,220} = 14.443, p < .001$), reflecting learning, and a group \times session interaction ($F_{11,220} = 2.178, p = .017$). This interaction appeared to be due to slower learning in the

ACCX group, which were impaired at the early stages of learning (simple effect of group significant for sessions 6,8,9 at $p < .01$) but reached the same difference score as shams by the end of session 12.

The impairment did not depend on the use of a difference score as the dependent measure, but was apparent when ratio scores (which are relatively independent of general activity levels) were analysed. Again, the ACCX group showed significantly poorer discrimination (effect of group: $F_{1,20} = 6.995$, $p = .016$). Assessed by this measure, the discrimination was poorer across all sessions (group \times session: $F < 1$, NS), though ratio scores increased during training (session: $F_{7,53,150.597} = 2.713$, $\tilde{\epsilon} = .685$, $p = .009$).




Figure 37. Two-stimulus discriminated approach task ($n = 12$ sham, 10 ACCX), showing impaired discrimination in the ACC-lesioned rats. These figures are based on the *number* of stimulus presentations during which a subject approached the food alcove, for direct comparison with the autoshaping task. Compare Figure 38. **A:** Raw approach scores, as a percentage of the total number of trials (which was 15). **B:** Difference scores (CS^+ approaches $- CS^-$ approaches). The maximum possible difference score is 15. **C:** Discrimination ratio scores (CS^+ approaches $\div (CS^+$ approaches $+ CS^-$ approaches)).

Analyses based on the proportion of time nosepeking to each stimulus

This approach score measures approach to a stimulus relative to that occurring during the VI. It proved a less sensitive measure than the number of trials on which approach occurred, as analysis of the proportion of time spent nosepeking only revealed an impairment in those animals with anterior cingulate lesions encompassing the ventral perigenual region (Figure 38).

The approach scores from all subjects were analysed using the model group₂ \times (session₁₂ \times stimulus₂ \times S). This showed a non-significant trend towards lower levels of stimulus-directed approach in the ACCX group (effect of group: $F_{1,20} = 3.574$, $p = .073$). There were main effects of stimulus ($F_{1,20} = 7.006$, $p = .015$) and session ($F_{5,876,117.53} = 48.928$, $\tilde{\epsilon} = .534$, $p < .001$), and a stimulus \times session interaction ($F_{7,235,114.701} = 2.78$, $\tilde{\epsilon} = .658$, $p = .009$), reflecting the acquisition of differential approach to the two stim-

uli. However, the stimulus \times group interaction did not reach significance ($F_{1,11} = 3.178, p = .09$), and no other terms involving group were significant ($F_s < 1, NS$). Interestingly, though, analysis of the sham and ACCX groups separately demonstrated significant stimulus discrimination in the shams (stimulus: $F_{1,11} = 13.487, p = .004$; session: $F_{4,176,45.936} = 24.84, \tilde{\epsilon} = .38, p < .001$; stimulus \times session: $F_{6,007,66.079} = 2.233, \tilde{\epsilon} = .546, p = .051$) but no evidence of discrimination in the ACCX group (stimulus: $F < 1, NS$; session: $F_{7,0,63,002} = 25.413, \tilde{\epsilon} = .636, p < .001$; stimulus \times session $F_{6,314,56.823} = 1.438, \tilde{\epsilon} = .574, NS$), despite similar group sizes (and therefore statistical power).

However, when the ACCX-whole subgroup were compared to shams, they were found to be significantly impaired. Despite the smaller number of animals, a stimulus \times group interaction was found ($F_{1,14} = 7.277, p = .017$), in addition to a main effect of session ($F_{5,402,75.621} = 30.581, \tilde{\epsilon} = .491, p < .001$) and a stimulus \times session interaction ($F_{7,155,100.172} = 2.105, \tilde{\epsilon} = .65, p = .048$). No other terms were significant ($F_s < 1.381, NS$). To explore the nature of the stimulus \times group interaction, data from each group were analysed using the model (session \times stimulus \times S). This demonstrated significant discrimination in the sham group, who approached more during the CS+ than during the CS- (stimulus: $F_{1,11} = 13.487, p = .004$; session: $F_{4,176,45.936} = 24.84, \tilde{\epsilon} = .38, p < .001$; stimulus \times session: $F_{6,007,66.079} = 2.233, \tilde{\epsilon} = .546, p = .051$), but no such discrimination in the ACCX-whole group (stimulus: $F_{1,3} = 1.312, NS$; session: $F_{3,594,10.783} = 12.534, \tilde{\epsilon} = .327, p = .001$; stimulus \times session: $F_{11,33} = 1.133, NS$).

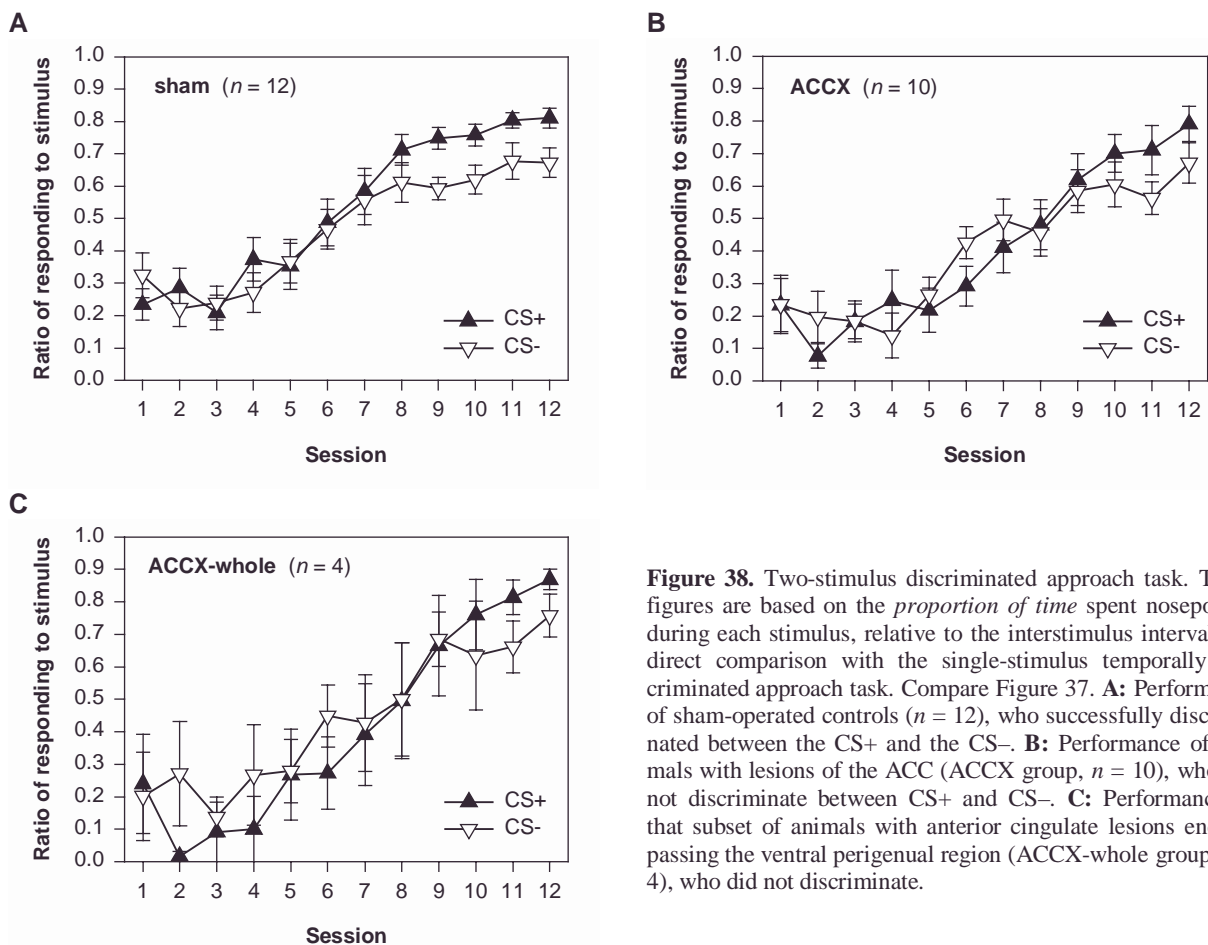


Figure 38. Two-stimulus discriminated approach task. These figures are based on the *proportion of time* spent nosepoking during each stimulus, relative to the interstimulus interval, for direct comparison with the single-stimulus temporally discriminated approach task. Compare Figure 37. **A:** Performance of sham-operated controls ($n = 12$), who successfully discriminated between the CS+ and the CS-. **B:** Performance of animals with lesions of the ACC (ACCX group, $n = 10$), who did not discriminate between CS+ and CS-. **C:** Performance of that subset of animals with anterior cingulate lesions encompassing the ventral perigenual region (ACCX-whole group, $n = 4$), who did not discriminate.




Figure 39. Two-stimulus conditioned reinforcement test. The figure shows the performance of sham animals, animals with lesions of the ACC, and the subgroup of those animals with lesions encompassing the ventral perigenual region of this cortex.

Two-stimulus conditioned reinforcement test

The sham-operated group preferred the CS+ to the CS− when allowed to respond for the two stimuli; thus, the CS+ served as a conditioned reinforcer. The ACCX group responded less, and showed poorer discrimination between CS+ and CS− (Figure 39). Square-root-transformed lever-press data were subjected to ANOVA using the model $group_2 \times (lever_2 \times session_2 \times S)$. Considering both groups together, there was a main effect of session ($F_{1,20} = 5.705, p = .027$), reflecting extinction. Subjects responded more on the CRf lever ($F_{1,20} = 10.832, p = .004$). There was also a session \times lever interaction ($F_{1,20} = 12.872, p = .002$); interestingly, this was due to improved discrimination on the *second* test day (simple effects analyses: effect of lever on day 1, $F_{1,20} = 3.476, p = .077$; effect of lever on day 2: $F_{1,20} = 19.484, p < .001$), which was due to a reduction responding on the NCRf lever but not on the CRf lever (orthogonal simple effects analyses: effect of session on CRf lever responding, $F < 1, NS$; on NCRf lever responding, $F_{1,20} = 18.276, p < .001$).

Animals in the ACCX group responded less on test (main effect of group: $F_{1,20} = 14.092, p = .001$). However, there were no other interactions involving group (group \times session: $F_{1,20} = 2.952, NS$; group \times lever: $F_{1,20} = 1.659, NS$; three-way interaction, $F_{1,20} = 1.615, NS$). However, it is clear from Figure 39 that discrimination was reduced in the ACCX group, and while the sham group on its own demonstrated significant discrimination between the levers (lever: $F_{1,11} = 9.117, p = .012$; lever \times session: $F_{1,11} = 11.138, p = .007$), in this analysis, the ACCX group did not (lever: $F_{1,9} = 2.72, p = .133$; lever \times session: $F_{1,9} = 3.085, p = .113$).

When the two sessions were considered separately for each group, the shams showed discrimination only on session 2 (simple effect of lever in session 1: $F_{1,11} = 3.315, NS$; in session 2, $F_{1,11} = 14.971, p = .003$). The ACCX group showed a similar pattern (simple effect of lever in session 1: $F < 1, NS$; in session 2, $F_{1,9} = 6.179, p = .035$). Thus, some discrimination was apparent in ACC-lesioned subjects, but it was much poorer than in sham-operated controls.

These conclusions were not materially altered by consideration of the ACCX-whole subgroup alone, except that these subjects showed a significant lever \times session interaction ($F_{1,3} = 10.668, p = .047$). While this might be interpreted as evidence of lever discrimination, Figure 39 shows that this was not the case: the interaction was due to a ‘crossover’, with the ACCX-whole subgroup responding more on the NCRf lever in session two. Considering each session separately, the ACCX-whole subgroup never showed discrimination (simple effect of lever in session 1: $F < 1, NS$; in session 2, $F < 1, NS$).

Incidentally, these results demonstrate in a within-subjects design that the CS+ was a more effective reinforcer than the CS− in sham-operated animals, eliminating a ‘stimulus-seeking’ explanation of their preference for the CRf lever in this task.

Summary

ACC-lesioned rats were significantly impaired at acquiring a discriminated approach response governed by two similar stimuli, only one of which was followed by reward. Like shams, they learned to approach during the CS+, but they also approached during the CS−, and exhibited much poorer CS+/CS− discrimination during acquisition. By at least some measures, they eventually acquired the discrimination, but took longer to learn it than shams. While the sham group responded more for the CS+ than the CS− in a test of conditioned reinforcement, the ACCX group responded less and did not discriminate to the same degree as shams.

Discussion

The results of this experiment provide clear support for Hypothesis 2: that the rat ACC contributes to discriminating similar stimuli on the basis of their association with reward, though it is not necessary for stimulus–reward associations *per se*. It is highly unlikely that the lesioned subjects simply failed to discriminate between the stimuli: ACC-lesioned rats have been shown to be normal (Bussey *et al.*, 1997b) or even improved (Bussey *et al.*, 1996) at tasks requiring left–right discrimination, and the stimuli used in the present task (and in the autoshaping experiments) differed in no way except their location. Similarly, ACC-lesioned rats have previously been shown to succeed in learning a CVD using stimuli with which they failed to learn an 8-pair concurrent discrimination (Bussey *et al.*, 1997b), again making a perceptual deficit an unlikely explanation. Nor is it plausible that a failure of response discrimination can account for the present results, as no response discrimination was required in the approach task (the responses measured following the CS+ and CS− were identical).

In the approach task, the CS+ and CS− may have served as instrumental S^Ds, just as in the one-stimulus version of the task used in Experiment 1. Indeed, it is not obvious that approach to a food alcove located *away* from the stimulus is in any sense a Pavlovian CR; thus, performance on the approach task may have been instrumental. Nevertheless, the CS+ predicted food delivery, so it was expected to enter into Pavlovian association with reward; in confirmation of this, the CS+ served as a CRf for both sham and ACC-lesioned rats, though discrimination was again much reduced in ACC-lesioned animals. Their poor discrimination is not simply attributable to generally low levels of operant responding (as ACC-lesioned rats acquired a free-operant response normally in Experiment 3), or failure to respond for conditioned reinforcement (given that they responded normally for CRf in Experiment 1). Thus, the failure of discrimination affected two kinds of measured behaviour, approach and instrumental responding.

GENERAL DISCUSSION

Contribution of the ACC to instrumental and Pavlovian behaviour

Lesions of the ACC have been shown to impair discrimination of reward- or punishment-associated stimuli in Pavlovian tasks, including autoshaping (present experiments; Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c) and autonomic conditioning (Powell *et al.*, 1994); in tasks whose Pavlovian/instrumental status is ambiguous (two-stimulus discriminated approach task, above; two-way active avoidance, Gabriel, 1993; 8-pair concurrent discrimination, Bussey *et al.*, 1997b); and in instrumental tasks depending on Pavlovian associations (responding for conditioned reinforcement, above). At present, the most parsimonious explanation is that the ACC forms or retrieves stimulus–outcome (Pavlovian) associations that may then influence instrumental behaviour, consistent with previous suggestions (Gabriel *et al.*, 1980a; Bussey *et al.*, 1996; Bussey *et al.*, 1997b).

In this section, the hypothesis will be developed that the ACC plays a specific role in *discriminating* stimuli on the basis of their association with reinforcement. However, there is also evidence that the ACC is particularly important *early* in learning (see also pp. 98–99); this will be considered first.

A time-limited role for the ACC?

The results of Experiment 4 also support the view of Gabriel and colleagues, derived from work with aversive conditioning in the rabbit, that the ACC has a time-limited role in learning (see pp. 98–99). Indeed, similar results have been obtained in PET studies of learning in humans (Raichle *et al.*, 1994; Petersen *et al.*, 1998). If the number of CS–US pairings is an important factor in the timecourse of the ACC's contribution to learning (and this may not be the case; see Poremba & Gabriel, 1999), an approximate quantification can be made from the present data. Judging by the difference scores from the two-stimulus approach task (Figure 37B, p. 109), the ACCX group reached the performance of the sham group at around session 10, after 150 CS⁺–US pairings. This exceeds the number of pairings that have been given in autoshaping acquisition tasks (Experiment 1; Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c), or autonomic conditioning experiments (Powell *et al.*, 1994), in which impairments have been observed. It corresponds roughly to the number of pairings experienced by the subjects in Experiment 2 (the autoshaping performance study) at the point when recovery was observed, as judged by the difference scores (though there was still a clear impairment in CS⁺ approach at this point). It would be extremely interesting, therefore, to investigate the effects of ACC lesions on autoshaping when made after considerable overtraining, say 250 or 500 CS–US pairings, as this might be beyond the point at which the ACC becomes unnecessary. At this point, behaviour would be expected to depend on other structures, such as the PCC (following Gabriel, 1993).

It is not easy to relate this quantitative estimate to the impairments found in operant discrimination tasks (Gabriel *et al.*, 1991a; Bussey *et al.*, 1996; Bussey *et al.*, 1997b), in which the impairments are typically measured by trials taken to reach a performance criterion. However, 150 US presentations is of the same order as the number of reinforcers required to establish an instrumental habit under a ratio schedule (Adams, 1982). It has been suggested that removal of the ACC leaves rats under the control of a S–R habit system (Bussey *et al.*, 1996), suggesting that experiments measuring the time after which the ACC is not required are actually measuring the speed at which a habit develops. It would be interesting to test an idea related to this hypothesis directly by administering an instrumental contingency test to ACC-lesioned rats (cf. Balleine & Dickinson, 1998a).

As discussed earlier (p. 99), ACC lesions impair autoshaping to a lesser degree when made following training than when made before acquisition (compare Experiment 4 to Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c), and some discrimination may eventually be attained by ACC-lesioned animals even if the lesions are made before acquisition (Parkinson *et al.*, 2000c, p. 49). This raises the possibility that the ACC is not required for *Pavlovian* responding after prolonged training. If this is indeed the case, what kind of representation controls behaviour thereafter? One highly speculative interpretation of the results summarized here is that Pavlovian skeletomotor conditioned responding can become ‘habitual’ through extended CS–US pairing, in the way that instrumental behaviour does. Psychologically, this would equate to Pavlovian responding being controlled by US-independent representations after prolonged training, more than it is after brief training (a suggestion testable by examining the effects of US devaluation on conditioned responding, as discussed in Chapter 1, pp. 21/25, and on pp. 98–99). Available behavioural evidence suggests the opposite (see Hendersen *et al.*, 1980; Mackintosh, 1983, p. 61) but the question has not been extensively studied.

Although conditioned responding may remain sensitive to US devaluation even after prolonged training, there is at least one Pavlovian process whose importance diminishes with overtraining. That is ‘mediated learning’, defined as the ability of a CS-activated representation of the US to enter into new associations (Holland, 1998). (Using a food US, mediated learning may be demonstrated by giving CS–US pairings, then pairing the CS with LiCl, and finally testing for an aversion conditioned to the US. Holland (1981) suggested that when the CS is paired with LiCl, the CS retrieves a representation of the US that can become associated with the LiCl, even though the food US is not physically present at that time.) Holland (1998) demonstrated that mediated learning occurred after brief amounts of initial training, but not after extended training, even though conditioned responding to the CS remained after extended training. On the basis of these results, Holland (1998) suggested that an *overtrained* CS maintains the ability to elicit a representation of the US for performance of the CR (‘mediated performance’), but this US representation cannot enter into new associations (cannot be used for mediated learning), either because it has reduced associability (following the theory of Pearce & Hall, 1980), or because mediated learning and mediated performance are embodied in distinct representational systems. Holland suggests one possibility of the latter kind: that mediated performance (simple conditioned responding) depends on the CS retrieving a *US-specific motivational value*, while mediated learning depends on the CS retrieving *specific sensory attributes of the US* (Holland, 1990b; 1990a; 1998). As discussed in Chapter 1 (p. 37) and by Holland (1998), there is strong evidence that the BLA is required for mediated performance and the retrieval of a US-specific motivational value, while it has been suggested that retrieval of US-specific sensory attributes depends on primary or higher-order sensory cortices, such as gustatory neocortex and rhinal cortex. It is not known whether the ACC is required for the transitory Pavlovian phenomenon of mediated learning. The present experiments and previous data, discussed above, suggest a transitory role for the ACC in learning. The present experiments also suggest a deficit in the CS specificity of Pavlovian associations in ACC-lesioned rats, but US specificity (which, according to this theory, is necessary for mediated learning) has not been tested, and it is not known how the ACC interacts with sensory cortex during learning. However, a deficit in mediated learning would not be sufficient explanation for the relative resilience of overtrained behaviour to ACC lesions observed by Gabriel and colleagues (see pp. 98–99).

The observation that ACC lesions do impair autoshaping performance (Experiment 2), even if transiently, suggests that the ACC does not merely ‘supervise’ learning in other systems, but stores or retrieves associations itself. The observation that the role of the ACC appears to diminish with time (reviewed above and on pp. 98–99) may be explained in two ways. As suggested above, the ACC might

contribute to a particular *form* of Pavlovian representation, whose importance normally diminishes with extended training. However, it is not necessary to postulate that the balance of representations controlling Pavlovian performance changes in order to explain the involvement of the ACC early in learning. The ACC might simply form a temporary store for the *same* kind of representations that eventually govern performance. It has been suggested that, early in training, associations are set up in the ACC and used to ‘teach’ other neural systems (Gabriel *et al.*, 1980a). The early and late representations might be of the same kind (be they CS–US_{sensory}, CS–US_{value}, CS–affect, etc.), with the ACC serving as a rapid but impermanent associator. On the basis of the evidence available to date, this seems the most likely explanation, but the issue is not settled.

Synthesis: a suggested role for the ACC in ‘disambiguating’ stimuli for its corticostriatal circuit

The ACC has been shown to be critical in a wide range of appetitive and aversive tasks in which two or more similar stimuli must be discriminated on the basis of their association with reinforcement (in the autoshaping and two-stimulus discriminated approach/conditioned reinforcement tasks presented in this chapter, and by Gabriel *et al.*, 1991a; Powell *et al.*, 1994; Bussey *et al.*, 1997a; Bussey *et al.*, 1997b; Parkinson *et al.*, 2000c). It is unlikely that these results reflect an attentional deficit (Muir *et al.*, 1996) or a failure of spatial discrimination (Gabriel *et al.*, 1991b; Powell *et al.*, 1994; Bussey *et al.*, 1996; Bussey *et al.*, 1997b). However, ACC-lesioned rats can discriminate between two stimuli of different modalities (Experiment 3, this chapter) and between two visual stimuli differing in a primary submodality such as colour (Bussey *et al.*, 1997b, Experiment 3). In at least some studies, ACC-lesioned animals have exhibited an early failure to discriminate between two CSs, but eventually improved or succeeded completely, implying that the early failure to discriminate was not due to a primary perceptual deficit (Experiment 4, this chapter; Gabriel, 1990; Gabriel *et al.*, 1991a; Gabriel, 1993; Parkinson *et al.*, 2000c). No deficits are apparent when ACC animals are required only to discriminate temporally between the presence and absence of a single CS, whether appetitive or aversive, and as judged by a wide variety of response systems; thus, ACC-lesioned rats were unimpaired at a single-stimulus discriminated approach task, responding for conditioned reinforcement, conditioned freezing, and PIT.

For at least one subset of these behaviours — locomotor approach — it seems very likely that the ACC influences behaviour through the Acb. The ACC projects strongly to the AcbC, which in turn projects to locomotor control regions of the ventral pallidum; lesions of the AcbC impair both autoshaping (Parkinson *et al.*, 2000c) and single-stimulus discriminated approach (Parkinson *et al.*, 1999b), and a functional connection between the ACC and the AcbC is necessary for autoshaping to develop (Parkinson *et al.*, 2000c). The effect of ACC and AcbC lesions on autoshaping differ, however; while ACC lesions typically result in ‘disinhibited’ responding to the CS– (Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c), AcbC lesions impair the conditioned approach response itself (Parkinson *et al.*, 2000c), just as they prevent conditioned approach to a single CS (Parkinson *et al.*, 1999b).

On the basis of these data, it is suggested that the ACC contributes to a sensorimotor aspect of conditioning (Parkinson *et al.*, 2000a). Without the ACC, animals can learn an ‘affective’ response to CSs; thus, they perform normally in the single-stimulus discriminated approach task, and exhibit PIT. They can also call up a motivational representation of the US (a role attributed to the BLA; Everitt *et al.*, 2000a), and so acquire a new response with conditioned reinforcement, and acquire conditioned freezing. However, CS specificity of these representations is impaired in ACC-lesioned rats; as a result, tasks that depend upon stimulus–reinforcer associations when similar stimuli must be discriminated require the ACC

(including autoshaping, and 8-pair concurrent visual discrimination). According to this hypothesis, the ACC disambiguates the stimuli for the rest of the limbic circuit of which it is a part (as illustrated in Figure 40).

For tasks in which the ventral striatum is the 'output' structure for behaviour, this 'extra' controlling circuitry may be a necessary refinement, as the striatum is itself anatomically capable only of discriminating amongst linearly separable cortical inputs (Wickens & Kötter, 1995, p. 206); on its own, the striatum should therefore be unable to perform an exclusive-or (XOR) discrimination (A^+ , B^+ , AB^-). Furthermore, discrimination of two linearly separable input patterns A and AB where $A \rightarrow \text{reward}$, $AB \rightarrow 0$ requires an inhibitory projection from unit B. As the direct cortical inputs to the striatum are all glutamatergic (excitatory), the striatum would seem unable to solve even this discrimination. However, the different cortical afferents to the Acb have been shown to gate each other's glutamatergic inputs (Cools *et al.*, 1991; Pennartz & Kitai, 1991; Floresco *et al.*, 1998) and in this sense, the ACC may operate to control the input of affective information (perhaps from the BLA) in order to direct motivational responses towards appropriate environmental stimuli. This hypothesis therefore predicts an impairment in configural or XOR discriminations in ACC-lesioned subjects.

Note that this account of ACC function does not suggest a primary sensory or perceptual role — ultimately, ACC-lesioned rats may make the sensory discrimination — but, more specifically, a role in the retrieval of appropriate affective information for specific stimuli that are attended to, and thus in the production of appropriate affective responses to stimuli (see also Turken & Swick, 1999). The concept that even early sensory representations may be neurally dissociated on the basis of the *response* for which the representation is used is not new (Goodale & Milner, 1992); from this perspective, the ACC may be critical for discriminating stimuli *for the purposes of stimulus–reinforcer associations*, but not for other perceptual processes. ACC-lesioned animals would be able to discriminate a CS+ from a CS– perceptually, but be unaware as to the correct stimulus towards which appropriate affective responses should be made.

This hypothesis can be shown to account for the impairment of avoidance learning by ACC lesions in rabbits (Gabriel, 1990; Gabriel *et al.*, 1991a; Gabriel *et al.*, 1991b), for in this task formation of specific stimulus–reinforcer associations confers an advantage. Indeed, in active avoidance behaviour an internally generated expectation of reinforcement may be particularly relevant, as successful behaviour results in the absence of primary reinforcement. As discussed earlier, the ACC is a site where discriminated activity (discharge to the CS+ but not the CS–) occurs early in discriminated avoidance training (Gabriel *et al.*, 1977). More generally, the ACC may provide stimulus–reinforcer information to other response systems. Thus, projections from the ACC to the CeA (see Fisk & Wyss, 1997), ultimately influencing brain-stem effector mechanisms, may direct autonomic responses toward appropriate environmental stimuli. This is supported by studies demonstrating a role for the ACC in the coordination of autonomic responses (Buchanan & Powell, 1982b; Neafsey *et al.*, 1993), and more directly by the finding that ACC lesions disrupt discriminated autonomic responses to a CS+ and CS– whilst not impairing the response itself (Powell *et al.*, 1994), much like the effects of ACC lesions on skeletomotor responses in the autoshaping procedure (Bussey *et al.*, 1997a; Parkinson *et al.*, 1999c). Finally, in tasks where stimulus–reinforcer learning is a disadvantageous strategy, ACC lesions can improve performance (Bussey *et al.*, 1996) (see p. 73).

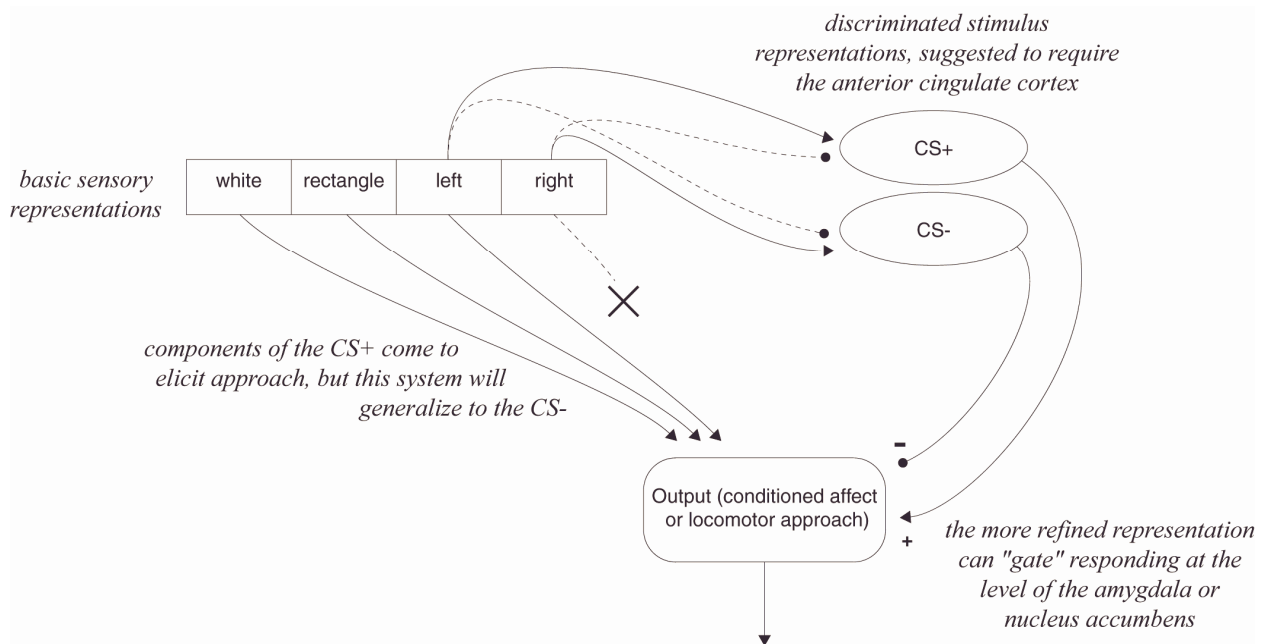


Figure 40 illustrates ‘disambiguation’ of stimuli, applied to autoshaping. In this example, the CS+ is a white rectangle on the left and the CS– is an identical stimulus on the right. Expression of autoshaping requires the CeA, AcbC, and Acb dopamine. In the absence of discriminated activity in the ACC, animals generalize from the CS+ to the CS–, impairing their behavioural discrimination in a disinhibited fashion. However, the animals will still discriminate between the presence and the absence of the CS+.

Interactions of the ACC with the amygdala and perirhinal cortex

The way in which the ACC interacts with the amygdala is far from clear, and requires further investigation. On the one hand, it has been suggested that discriminated neuronal activity in the ACC depends on the amygdala. Poremba & Gabriel (1999), using rabbits, found that inactivation of the amygdala using the GABA receptor agonist muscimol blocked discriminated avoidance learning and cingulate neuronal plasticity, replicating previous findings with electrolytic lesions (Poremba & Gabriel, 1997). Furthermore, the amygdala-inactivated rabbits showed no savings effect, suggesting that they failed to learn while the amygdala was inactivated, not that they simply failed to perform the response. However, amygdala inactivation had no effect on well-trained animals. (Indeed, further training was not necessary to reduce the amygdalar involvement; simple passage of time was enough!)

On the other hand, studies of autoshaping do not suggest the involvement of the BLA in ACC-dependent tasks: BLA-lesioned rats have been shown to acquire normal discriminated autoshaping behaviour (Parkinson *et al.*, 2000b). Poremba & Gabriel (1999) suggested that the amygdala is critical for learning about ‘emergency’ situations involving stimuli of extreme motivational valence, which might account for the studies just described, but the BLA is clearly involved in other appetitive tasks (Everitt *et al.*, 1999; Everitt *et al.*, 2000a).

It is possible that these data may be reconciled by closer consideration of the anatomical site concerned. Poremba & Gabriel (1999) aimed their muscimol injections at the BLA, but did not specify the degree to which the CeA was affected. In their earlier study (Poremba & Gabriel, 1997), the CeA was damaged by the electrolytic lesions used, and there was a significant correlation between CeA (and lateral amygdala) damage and performance of the avoidance task, while this correlation was not significant for the BLA. These results are easily reconciled with autoshaping studies showing that the CeA, but not the BLA, is critical for the development of autoshaping (Parkinson *et al.*, 2000b). Given that the amygdala

plays a time-limited role in the acquisition of the avoidance task used by Gabriel and colleagues (Poremba & Gabriel, 1997), it is extremely interesting to note that lesions of the CeA do *not* impair the performance of a well-trained autoshaped response (Everitt *et al.*, 2000b) at a stage of training when ACC lesions do (Experiment 2). These results would be compatible with a role for the CeA in the learning process through which the ACC acquires specific stimulus–reinforcer associations in this task. As the amygdalocortical projections to the ACC arise predominantly from the basal nucleus (see Amaral *et al.*, 1992, pp. 46–47), it may be that the CeA influences cortex through its projections to the chemically defined systems of the brainstem reticular formation. There are precedents for this suggestion (see Everitt *et al.*, 2000a); for example, the CeA has been shown to have a role in upregulating the associability of conditioned stimuli via its projections to the basal forebrain cholinergic systems (Holland & Gallagher, 1993a; Chiba *et al.*, 1995; Holland, 1997), which project to wide areas of cortex, including the ACC (Butcher, 1995); in turn, regulation of cortical CS representations has been shown to depend upon the cholinergic innervation of the cortex (Weinberger, 1995).

Finally, while the ACC has been implicated in stimulus–reinforcer learning, investigations of stimulus–*stimulus* learning have implicated the perirhinal cortex (PRh) as a critical site for complex, cross-modal and configural associations (e.g. Saksida & Bussey, 1998; Murray & Bussey, 1999; Murray *et al.*, 2000; Nicholson & Freeman, 2000; Saksida *et al.*, 2000). It remains to be established whether the ACC and PRh interact when complex stimuli are associated with reinforcement, and what their relative contributions to behaviour are.

Comparison with other interventional studies in rodents

Can the hypothesis of ACC function developed above explain results from studies using very different paradigms?

Mice

Meunier *et al.* (1991) have studied spatial discrimination learning and reversal in a T-maze using mice. Mice with ACC lesions learned the initial acquisition and first reversal at the normal rate, but they were impaired during subsequent reversal sessions, failing to show positive learning transfer when compared to controls. Yet when all they had to do was learn the *same* discrimination over several days (repetition), there was no impairment. In fact, they learned the first repetition more easily than control animals. Meunier *et al.* interpreted the ACC deficit as an inability to remember the temporal order of previously acquired spatial responses, though the integrity of each individual response was maintained.

It is worth noting the timescale implied by this hypothesis. The reversal sessions of Meunier *et al.* (1991) were on consecutive days (in any one session, one arm of the T-maze was baited consistently). A memory for temporal order only helps in the solution of the reversal task if the animal is following a rule of the sort: ‘what was right yesterday? Let me do the other today.’ As justification for their hypothesis, Meunier *et al.* (1986) have shown that mice with ACC lesions can still demonstrate interproblem transfer in a maze task supposed to require the formation of a general rule, but not a rule involving temporal order. But this is insufficient justification to call this phenomenon ‘memory for temporal order’. The general form of the ‘temporal order’ hypothesis also makes clear predictions; for example, rats with ACC lesions should be impaired on a discrimination task using three stimuli presented in a sequence, where ABC→reward and BAC→no reward; this has not yet been tested.

If these findings are re-examined in the light of Bussey’s (1996) hypothesis, it could be that the ACC-lesioned mice in Meunier’s (1991) study successfully learned S–R associations that allowed them to perform normally on the initial acquisition session, first reversal and repetitive tests. PCC lesions impaired

mice on exactly these tests. However, the presence of the ACC might confer an ability to respond rapidly and flexibly to an environment with changing stimulus–reinforcement relationships, withholding responses to unrewarded stimuli, which in normal mice results in improved performance over the course of reversal training.

Rats

Interventional studies using rats have revealed other features of the phenotype of ACC lesions that are not all easy to encompass within the hypothesis outlined above. In particular, they emphasize *disinhibition* and *over-responding* in ACC-lesioned rats. Weissenborn *et al.* (1997) studied the acquisition of responding for intravenous cocaine under second-order schedules of reinforcement. ACC-lesioned rats exhibited greater locomotor activity (both spontaneous and cocaine-induced), they were more likely to self-administer excessive amounts of cocaine during acquisition, and while their dose–response curve was normal on a FR1 schedule, they responded at high rates throughout the fixed-interval phase of the second-order schedule, exhibiting an attenuated fixed-interval ‘scallop’. Weissenborn *et al.* related this to Bussey’s (1997a) hypothesis by suggesting that the rats had failed to learn the significance of the cocaine-associated stimulus that normally maintains responding on this schedule. Such hyperactivity was not found in the present series of experiments; as discussed on pp. 90/97, this may have resulted from differences in lesion site. Another factor to be considered in Weissenborn *et al.*’s (1997) experiments was chronic cocaine experience, which might interact with the effects of ACC lesions.

Muir *et al.* (1996) studied a five-choice serial reaction time task (5CSRTT) in which rats must wait for the presentation of one of five brief visual stimuli, and then respond at the location of the stimulus in order to gain reward. Muir *et al.* found that ACC lesions had no effect on the accuracy of visual attentional performance, either at baseline or with superimposed attentional manipulations (varying the stimulus duration or the ITI, or interpolating bursts of white noise). However, the lesions increased the number of premature, anticipatory responses (in which the animal responds before a stimulus has been presented), increased the number of ‘perseverative’ responses (in which the animal responds several times to the location where a stimulus was recently presented), and decreased the number of errors of omission. The same animals performed normally on a test of passive avoidance, in which electric shock is delivered in one half of a two-chamber apparatus and the subject subsequently avoids the ‘dangerous’ chamber.

Clearly, these results may be explained in terms of disinhibited or impulsive motor responding. The results of Muir *et al.* (1996) suggest that the ACC-lesioned rats were unable to withhold responding to locations where rewarded stimuli were intermittently presented. However, there was no evidence of such a deficit in the present series of experiments; locomotor hyperactivity was not apparent, no test of free-operand responding demonstrated hyperactivity, and ACC-lesioned rats did not over-respond to the location of a rewarded CS when that CS was not present. Differences in lesion site may partly be responsible for these discrepancies — the present lesions were more anterior to those used by Muir *et al.* (1996, Figure 1C) (see also Figure 14 caption, p. 72) and recent results suggest that ACC lesions centred on the perigenual region, similar to those used in the present experiments, do not produce deficits on the 5CSRTT (A. Christakou, unpublished observations; personal communication, 10 October 2000). It should be noted that the psychological basis of premature responding in the 5CSRTT is not well understood; however, it is not clear how these results can be reconciled in terms of a single deficit. Investigating the role of the ACC in explicit tests of ‘motor impulsivity’ (see Evenden, 1999b) such as the go/no-go task (e.g. Harrison *et al.*, 1999) may be well worth while. This task (in which subjects must respond to one stimulus but withhold responding to another stimulus) has the added advantage that the degree to which one response is prepotent can be varied by altering the relative proportion of ‘go’ and ‘no-go’ trials, pro-

viding a further test of inhibitory control. The role of the ACC in impulsive *choice* will be investigated in Chapter 7.

Other studies of the rat ACC have frequently concentrated on the region directly superior to PL (equivalent to dorsal mPFC in Figure 14C, p. 72), an area that was not the focus of the present experiments. Despite the differences in location, there are some commonalities among findings. For example, such lesions have minimal effects on rats' spatial discrimination or working memory (Neave *et al.*, 1994; Ragozzino *et al.*, 1998) or their ability to switch strategies between the use of visual and spatial cues (Ragozzino *et al.*, 1999), yet produce severe impairments in a number of radial maze tasks (Seamans *et al.*, 1995): rats with reversible (lidocaine) lesions of the ACC preferentially revisit previously baited arms. This last deficit has clear analogies with the 'disinhibited', perseverative behaviour observed in the 5CSRTT by Muir *et al.* (1996), but might also be explicable in terms of a failure to inhibit responding to unrewarded stimuli (maze arms) in a situation in which there are many stimuli, differentially associated with reward, and in which the rewarded stimulus changes rapidly. Seamans *et al.* (1995, p. 1071) describe the ACC as providing response flexibility by suppressing the effect of simple stimulus–reward associations on behaviour, an interpretation clearly compatible with the present results.


Homology between rodent and primate ACC

In order to examine the present results in the context of primate lesion and imaging studies, it will be vital to consider the homology between rodent and primate ACC. The various terminologies used to describe rat ACC are summarized by Neafsey *et al.* (1993) and reproduced in Table 11. The lesions used in the present series of experiments were of Cg1 and Cg2 situated at, and caudal to, the genu of the corpus callosum, corresponding to area 24a and caudal area 24b in the rat (see Figure 14, p. 72; Table 9, p. 71; and Table 11). In turn, these areas have a homologue in monkey and human ACC, as judged by their pattern of afferent and efferent connections (Öngür & Price, 2000); Figure 41 depicts rat prefrontal cortex and maps of macaque and human PFC that were designed to represent homologous regions with the same area number (Öngür & Price, 2000). A rough equivalence may therefore be drawn across the three species.


Table 11. Cytoarchitectonic subdivisions of anterior cingulate cortex, from Neafsey (1993); compare Figure 14 (p. 72). (Abbreviations are the same as those in Figure 41, with the addition of Prcm, medial precentral cortex; Cg1–Cg3, cingulate cortex; HP, hippocampal rudiment.)

Source	Dorsal						Ventral
Krettek & Price (1977)	Prcm	ACd	ACv	PL	IL	DPC	TT
Krieg (1946); Vogt & Peters (1981)	4	24b	24a	32	25	25	TT
Zilles & Wree (1985)	Fr2	Cg1	Cg2	Cg3	IL	IL	HP
Uyling & van Eden (1990)	Fr2	ACd	ACv	PL	IL	IL	TT

A. Rat




B. Macaque monkey (1)




C. Macaque monkey (2)

A. CYTOARCHITECTURE



B. FUNCTIONAL SUBDIVISIONS



D. Human




Figure 41. Medial prefrontal cortex in rats, monkeys, and humans (not to the same scale). The top of the diagram is the superior direction in each panel. **A:** Medial frontal cortical regions in the rat, rostral to the right, from Neafsey *et al.* (1993); compare Figure 14 (p. 72). The number-based designations from Table 11 (p. 120) have been superimposed upon the original figure. (Fr2, frontal cortex 2; ACd/ACv, dorsal/ventral anterior cingulate cortex; PLd/PLv, dorsal/ventral prelimbic cortex; IL, infralimbic cortex; MO, medial orbital cortex; DPC, dorsal peduncular cortex; TT, taenia tecta; OB, olfactory bulb.) The corpus callosum is shown in black. **B:** Orbital and medial prefrontal cortex in the macaque monkey, rostral to the left, from Carmichael & Price (1994) via Öngür & Price (2000) (cc, corpus callosum). **C:** Rhesus macaque monkey cingulate cortex, rostral to the left, from Vogt *et al.* (1992), showing functionally specialized regions. (CS, cingulate sulcus; VMA, visceromotor area; VOA, vocalization area; NCA, nociceptive area; CMAr/CMAv, rostral/ventral cingulate motor areas; AAA, attention-to-action area; VSA, visuospatial area.) **D:** Human medial prefrontal cortex, rostral to the right, from Öngür & Price (2000).

Interventional studies in primates

Most studies of the primate ACC are rather unhelpful for comparison to the present studies of the contribution of the ACC to conditioned behaviour. This is for two reasons. Firstly, many primate interventional studies concerned with the ACC have used non-excitotoxic lesion techniques (see Devinsky *et al.*, 1995). The nature of the lesion is critically important in this region; any lesion that destroys the cingulum bundle will disconnect large portions of cortex, for this bundle contains not only all afferent and efferent connections of the cingulate cortex, but also fibres that pass to and from the rest of the prefrontal (including orbitofrontal) cortex, notably the reciprocal connections between the prefrontal cortex and the medial temporal lobe (Vogt, 1993, pp. 24–25). (The functions of the orbitofrontal cortex were briefly reviewed in several contexts in Chapter 1, pp. 39/45/55, and will not be considered here.) Secondly, most studies have concentrated on unconditioned behaviour. It is clear that regions of the primate ACC are involved in a bewildering range of motivationally-oriented unconditioned behaviour. Devinsky *et al.* (1995), reviewing these studies, consider the ACC to be the ‘anterior executive’ region of the cingulate cortex (cf. Figure 41C); the ACC is further subdivided into an ‘affect’ region (area 25 and rostral area 24) and a ‘cognition’ region (caudal areas 24’ and 32’, nociceptive cortex and the cingulate motor areas).

Functional subdivisions of the rhesus monkey ACC are shown in Figure 41C. As the regional names would suggest, the primate ACC has been implicated in the perception of pain (and by reference to the rodent literature, avoidance learning); as a part of premotor cortex; in visceromotor control (and, again by reference to the rodent literature, classically conditioned autonomic responses); and in vocalization that has social or emotional content (Devinsky *et al.*, 1995), interpretations that are supported by stimulation studies in humans and other primates. In addition, the ACC has been implicated in action selection or ‘attention to action’ (discussed below). This last concept has the most relevance to the present rat experiments, and has been best studied in the human; therefore, studies of the human ACC will be considered next.

Correlational studies in humans

Isolated destruction of the human ACC is rare (Devinsky *et al.*, 1995), so lesion studies of humans have mostly been of patients with frontal lobe tumours. Lesions of the ACC have produced a wide variety of symptoms, including apathy, inattention, autonomic dysregulation, emotional instability, and akinetic mutism (Devinsky *et al.*, 1995; Bush *et al.*, 2000). However, such studies are often compromised by a lack of anatomical specificity: tumours and epileptic foci do not respect anatomical boundaries, and if these tumours involve the ACC, their resection inevitably compromises the cingulum bundle, and thus orbitofrontal cortex function. Indeed, many of the patients studied by Damasio and colleagues (see Chapter 1, p. 55) have ACC damage in addition to orbitofrontal lesions (Bechara *et al.*, 2000). Some of the best studies of human ACC are therefore correlational, in that they aim to observe differences in ACC activity that are correlated with task performance or mental state, without using interventional techniques to alter ACC function and observe the effect on behaviour. While interventional techniques are required to show that the ACC has a causal role in a particular aspect of behaviour, correlational techniques have provided useful information about ACC function.

Emotional states, emotionally significant stimuli, and mood

The anterior, ventral ACC (Brodmann’s areas 24a/b and 25), part of the ‘affective’ subdivision of the ACC (Devinsky *et al.*, 1995), is now strongly implicated in the pathology of depression in humans (Bench *et al.*, 1992), as well as in the control of normal mood. Drevets *et al.* (1997) observed that this

area of the ACC ('subgenual prefrontal cortex' or subgenual area 24; see Öngür *et al.*, 1998) showed decreased blood flow in unmedicated familial bipolar and unipolar depressives using positron emission tomography (PET), though this was in part due to a reduced grey matter volume as assessed by magnetic resonance imaging (MRI); if this is corrected for, blood flow per unit volume was increased (Mayberg, 1997; Drevets, 2000). Mayberg *et al.* (1994; 1996; Mayberg, 1997) have demonstrated similar abnormalities; metabolic activity in rostral ACC (rostral area 24a/b) is also unique in differentiating those depressed patients who eventually respond to pharmacological antidepressant therapy from those that do not (Mayberg *et al.*, 1997). Areas 24a/b and 25 are also part of a cortical network whose metabolic activity alters in normal sadness (Mayberg *et al.*, 1999). Mayberg *et al.* (1999; Mayberg, 2000), reviewing these data, have suggested that hyperactivity of subgenual area 24/area 25 is a primary factor in sadness and depression, causing reciprocal suppression of metabolism in adjacent ACC and dorsolateral prefrontal cortex, which may explain the efficacy of surgical destruction of the subgenual cingulate as a therapy for refractory depression.

Imaging studies have also shown that the human ACC responds to emotionally significant stimuli. It is reliably activated by cocaine-associated cues in cocaine users, more so than by neutral stimuli in the same individuals, or by cocaine-associated cues in non-users (Maas *et al.*, 1998; Childress *et al.*, 1999; Garavan *et al.*, 2000); such activation may be associated with cocaine craving (e.g. Volkow *et al.*, 1996; Volkow *et al.*, 1997; Maas *et al.*, 1998; Childress *et al.*, 1999). While fewer studies have examined the effects of natural reinforcers, it appears that the ACC is similarly activated by emotionally significant non-drug stimuli in normal humans (sexual images; Garavan *et al.*, 2000).

Attention, conflict monitoring, error detection, and action selection

Attention and action

In humans, PET studies have provided evidence that the ACC is involved in executive attention. In attentional target detection tasks, blood flow increases with the number of targets to be detected, while flow to the anterior cingulate gyrus is reduced below baseline during the maintenance of vigilance (reviewed by Posner, 1995, pp. 620–621). These PET studies have also suggested a role for the ACC in 'willed' tasks, perhaps with a motivational role; along with dorsolateral PFC, blood flow to ACC is significantly increased in tasks requiring a voluntary choice of action, compared to routine, well-rehearsed actions (Frith *et al.*, 1991).

Detecting errors: the error-related negativity (ERN) and its localization to the ACC

An event-related brain potential (ERP) is an electroencephalographic (EEG) potential that has been time-locked to an event. While studying choice reaction times (RTs) in humans, it was observed that a negative EEG potential was evoked when subjects made an error (Falkenstein *et al.*, 1990; Gehring *et al.*, 1990; Gehring *et al.*, 1993). This potential was named the error-related negativity (ERN).

The literature on the ERN is large and will be summarized briefly (for reviews, see Brown, 1999; Falkenstein *et al.*, 2000; Scheffers & Coles, 2000). The ERN begins to develop at around the time of the erroneous response, and peaks ~100 ms later; it is small or non-existent following correct responses. The ERN is hypothesized to reflect part of a process in the brain that monitors ongoing actions, compares them with intended actions, detects any mismatch, flags the presence of an error if mismatch exists, and takes action to correct ongoing or future performance (e.g. Gehring *et al.*, 1993; Bernstein *et al.*, 1995; Miltner *et al.*, 1997). There is wide support for this general view. ERNs are generated in a variety of tasks, including visual and auditory discriminations, go/no-go tasks, and the Eriksen flankers task, in

which subjects must respond to the identity of a briefly-presented target letter (H or S) that is either flanked by compatible letters (e.g. HHHHH) or by letters associated with the alternative response (e.g. SSHSS) (Gehring *et al.*, 1993). This generality suggests that the ERN is not a reflection of the stimuli used in the tasks. The ERN also occurs regardless of the particular motor response being measured (Holroyd *et al.*, 1998). Nor does it depend on a particular type of error: in one go/no-go task, subjects may make errors of choice, in which they respond with the incorrect hand on ‘go’ trials, or errors of action, in which they respond (with either hand) on ‘no go’ trials; both these conditions generate an ERN (Scheffers *et al.*, 1996). The error-detection process appears to rely on the same representations that govern task performance: for example, when well-practised subjects perform RT tasks for long periods without sleep, their performance worsens as a result of impaired perceptual processing, and the ERN declines as the representation of the correct response is degraded (Scheffers *et al.*, 1999). The ERN also reflects subjects’ perception of accuracy. Scheffers & Coles (2000) gave subjects a task in which the visual stimuli governing performance were degraded. Regardless of behavioural accuracy, the ERN at the time of responding correlated with subjects’ subsequent reports of how inaccurate the response was — that is, errors perceived as such were associated with large ERNs, but so were correct responses perceived as errors. The ERN was smaller on trials where the subject was unsure whether an error had been made (due to limitations on the available data), and smallest when the subject thought he had responded correctly (even when an error had been committed). Finally, when the subject must learn to respond based on a delayed feedback signal, an ERN is generated in response to feedback indicating incorrect performance (Miltner *et al.*, 1997), a moment at which no response is being made.

Much of the controversy about the precise significance of the ERN is attributable to the bidirectional hypothesis of its function stated above: errors are suggested to generate the ERN, and the ERN is suggested to correct errors. Thus, the ERN is associated with conditions of error; for example, it is larger when the task instructions emphasize response accuracy over speed (on trials matched for RT; Gehring *et al.*, 1993) and when responses are late in a task in which speed is emphasized (Luu *et al.*, 2000b); greater motor discrepancies between intended and actual responses also generate larger ERNs (Bernstein *et al.*, 1995). However, the ERN is also associated with correction processes: large ERNs are also associated with less forceful errors that are more likely to be followed by correction responses (with longer RTs), and large ERNs are associated with more conservative behaviour in the future (see Scheffers & Coles, 2000).

In support of early speculations (Gehring *et al.*, 1993), equivalent dipole analyses, together with neurophysiological and biophysical considerations, point to the ACC as the likely source of the ERN (Dehaene *et al.*, 1994; Coles *et al.*, 1998; Bush *et al.*, 2000) — indeed, the ERN may have first been noticed by researchers recording directly from the ACC (area 24) in macaque monkeys (Gemba *et al.*, 1986). The ACC has thus been likened to a supervisory attentional system (Norman & Shallice, 1986) (see Grossman *et al.*, 1992). Given the importance of error signals in many models of learning (famously, that of Rescorla & Wagner, 1972), there has been considerable interest in relating the ERN to learning (see Kopp & Wolff, 2000; Schultz & Dickinson, 2000). Another well-studied candidate for an error signal in the brain is the dopamine system; midbrain dopamine neurons in the SNc/VTA fire in response to primary rewards, but come to respond instead to signals predictive of reward, and signal discrepancies between predicted and experienced rewards (Mirenowicz & Schultz, 1994; Schultz *et al.*, 1995b; Mirenowicz & Schultz, 1996; Schultz *et al.*, 1997). This has led to the incorporation of the DA signal in models of learning, most notably those based on the algorithm entitled *temporal difference* (TD) learning (Sutton, 1988; Barto, 1995; Houk *et al.*, 1995). In an intriguing development, it has been suggested that

the TD error signal conveyed by DA neurons is responsible for the ERN in the ACC (Holroyd *et al.*, 1999) — intriguing not least in relation to the suggestion (p. 117) that the CeA, a likely controller of the VTA, may regulate ACC function. However, the data summarized here suggest that the ACC's functions are more to do with response errors than errors of reward prediction (Schultz & Dickinson, 2000).

Finally, the ERN has been shown to be abnormal in psychopathological states to which the ACC has been suggested to contribute. A potential link from the ERN literature to the involvement of the ACC in depression (discussed above, p. 122) has been provided by Luu *et al.* (2000a), who found that the ERN was larger in normal humans who scored highly for the personality dimensions of 'negative emotionality' and 'negative affect' as assessed by questionnaires. Similarly, Gehring *et al.* (2000) found that the ERN was larger in patients suffering from obsessive-compulsive disorder (OCD), a disorder in which self-monitoring and error correction may be pathologically enhanced, in which ACC metabolism is abnormal, and which has been successfully treated surgically by cingulotomy (Baer *et al.*, 1995; Devinsky *et al.*, 1995; Gelder *et al.*, 1995, pp. 180–185; Breiter *et al.*, 1996; Busatto *et al.*, 2000).

Activating the ACC: response competition and the Stroop test

Studies of the ERN are supported by an array of functional imaging data implicating the ACC in error-related tasks. The spatial precision of PET and functional MRI (fMRI) far exceeds that of the EEG; thus, the anatomical basis of activation focus can be accurately localized. However, the temporal resolution of PET and fMRI is far poorer than the EEG (while the technique of magnetoencephalography or MEG, which has high spatial and temporal resolution, is presently only suitable for superficial cortical sites). Inevitably, the poorer temporal resolution of functional imaging has led to controversy about the significance of activation foci within the ACC.

The Stroop test (Stroop, 1935) is a prototype of the tasks that increase metabolic activity in the ACC. In the Colour Stroop (Figure 42), the subject must report the colour of a series of words, while ignoring the word itself. The task has congruent and neutral conditions, in which the word itself helps or does not contribute to performance (Figure 42, left and middle columns), but in the incongruent condition, each word is the name of a colour that differs from the colour in which the word is printed (Figure 42, right). The incongruent condition of many variants of the Stroop test strongly activates a focus in the ACC (Pardo *et al.*, 1990; Bench *et al.*, 1993; Carter *et al.*, 1995; Derbyshire *et al.*, 1998; MacDonald *et al.*, 2000), just as the Stroop test elicits an ERN from the ACC (Liotti *et al.*, 2000). The precise locus depends on the nature of the task; thus, while the Counting Stroop (prototype: report the number of words present, even when the words are numbers) activates the 'cognitive', caudal division of the ACC, the Emotional Counting Stroop (prototype: count neutral or emotionally-charged words, such as MURDER) activates the 'affective' division, rostral and inferior to the genu of the corpus callosum (Bush *et al.*, 1998; Whalen *et al.*, 1998; Bush *et al.*, 2000; MacLeod & MacDonald, 2000).

Evaluation of errors, or evaluation of response competition (conflict)? As the incongruent condition of the Stroop test activates the ACC even when behavioural performance is accurate (see MacLeod & MacDonald, 2000), it has been suggested that the ACC evaluates the degree of response competition or conflict, rather than simply detecting errors. In a different continuous performance task, Carter *et al.* (1998) similarly observed that the ACC is not only activated when incorrect responses are made, but when correct responses are made under situations of high response competition. Similar results have been obtained by Rogers *et al.* (1999), who observed ACC activation that was correlated with response conflict using a decision-making task in which error rates were held constant. Carter *et al.* (1998; 1999) suggest that rather than implement a comparator process (correct versus actual responses), the ACC monitors compe-

blue	willow	red
yellow	trek	green
red	armchair	blue
green	prefect	yellow
blue	felicitous	blue
green	destructive	green
yellow	milk	yellow
blue	bore	blue
red	selection	red
yellow	karyotype	green

Figure 42. Activate your ACC: a version of the Stroop test. The subject is asked to read aloud the colour each word is printed in (ignoring the word itself), as accurately and rapidly as possible. The left-hand column illustrates a congruent condition, the middle column is a neutral condition, and the right-hand column is an incongruent condition. There is a reaction time cost for the incongruent condition and this condition strongly increases metabolic activity in the ACC.

tion between responses. They argue that ‘conflict monitoring’ is a better description of ACC function than ‘error detection’.

Evaluation, or strategic selection of actions? Investigators have also sought to define whether the ACC is primarily evaluative, detecting errors or response conflict, or strategic, implementing ‘selection for action’ (a term originated by Allport, 1987, defined as ‘processes that reduce the competition between potential responses to a stimulus’) (see Carter *et al.*, 2000). As Devinsky and colleagues state, ‘when a response selection is made, including the decision not to move, area 24’ is engaged’ (Devinsky *et al.*, 1995, p. 298). There is experimental evidence both for an evaluative interpretation (e.g. Botvinick *et al.*, 1999; Carter *et al.*, 2000) and a strategic interpretation (e.g. Paus *et al.*, 1993; Awh & Gehring, 1999; Turken & Swick, 1999). When examining action selection, dissociations within the ACC have also been observed for different response modalities, suggesting that the ‘executive control’ functions of the ACC are separable according to the response system being controlled (Paus *et al.*, 1993; Awh & Gehring, 1999; Turken & Swick, 1999).

Two criticisms can be levelled at this approach, one practical and one functional. The practical problem with this approach is the potential for a bidirectional relationship between errors and ACC activation (for example, one might expect the following sequence: more errors → ACC activation → correction → fewer errors). This possibility complicates many of the studies cited (see MacLeod & MacDonald, 2000), especially when one allows that error correction may occur before the action is made. For example, an error-detector might be involved in the incongruent Stroop test because subjects start to generate internal representations of multiple responses (to the word and to the colour), one of which triggers an internal error signal, leading to on-line correction of behaviour. Response competition and error detection may share features. Part of the reason for the success of ERN research is that the EEG technique allows trial-by-trial monitoring (something that is difficult using PET or fMRI), but part has been due to experimental technique that breaks the bidirectionality described above — for example, by providing an error-related signal in the absence of responding (Miltner *et al.*, 1997), or by measuring the ERN when subjects’ belief about the accuracy of their responses differs from the actual accuracy (Scheffers & Coles, 2000).

On a functional level, the distinction between evaluative and strategic functions may be — at least in part — doomed to failure. If an error-detecting process cannot correct errors, what good is it? If a super-

visory action-selection mechanism is not activated when response competition or the likelihood of error is high, then when? In this respect, consideration of the interaction *between* structures is likely to be as helpful as the consideration of the structures themselves. Recent studies are beginning to take this approach, considering, for example, the contribution of the dorsolateral PFC to the function of the human ACC (Cohen *et al.*, 2000; Gehring & Knight, 2000).

Relating human and rodent studies

It would be optimistic to be able to relate the entire literature on human ACC function to studies of rats, mice and rabbits. In particular, there is little evidence to address the question of whether the rodent ACC responds to errors or response-conflict situations (though the macaque ACC does; Gemba *et al.*, 1986), and there are few anatomically well-specified human lesion studies investigating the behavioural role of the ACC. However, common themes can be drawn. The rostral division of the human ACC responds to stimuli of affective significance (e.g. Whalen *et al.*, 1998), as does the rabbit ACC (Gabriel *et al.*, 1980a; Gabriel *et al.*, 1980b; Gabriel & Orona, 1982; Gabriel *et al.*, 1991b). The rabbit ACC uses this information to contribute to the selection of actions in instrumental avoidance tasks, a function similar to that attributed to the human ACC, and both the human and the rodent ACC control a wide variety of skeleto-motor and autonomic response systems (e.g. Paus *et al.*, 1993; Powell *et al.*, 1994; Devinsky *et al.*, 1995; Bussey *et al.*, 1997a; Awh & Gehring, 1999; Turken & Swick, 1999). The rat ACC contributes to the control of behaviour when faced with two or more similar stimuli predicting different outcomes (present experiments, and Gabriel *et al.*, 1991a; Powell *et al.*, 1994; Bussey *et al.*, 1997a; Bussey *et al.*, 1997b; Parkinson *et al.*, 2000c); analogies may be drawn with human 'response conflict' accounts. The human ACC is suggested to be activated by novelty or errors (Falkenstein *et al.*, 1990; Gehring *et al.*, 1993; Dehaene *et al.*, 1994; Berns *et al.*, 1997; Coles *et al.*, 1998) and thus to be involved in learning (Kopp & Wolff, 2000; Schultz & Dickinson, 2000); it is activated early in the acquisition of new tasks (Raichle *et al.*, 1994; Petersen *et al.*, 1998). Similarly, the contribution of rodent ACC is most marked early in training, when most learning might be expected to occur (see pp. 99/113); the monkey ACC ERN is present only during learning, when errors are still being made (Gemba *et al.*, 1986), and the mouse ACC appears to contribute to performance when response–outcome contingencies are changing rapidly (Meunier *et al.*, 1991). It is to be hoped that future studies will begin to bridge these two literatures.

Chapter 4.

Role of the nucleus accumbens core and shell in Pavlovian–instrumental transfer

Abstract. When an initially neutral stimulus has been paired in Pavlovian fashion with an appetitive outcome such as food, noncontingent presentation of this conditioned stimulus (CS) can enhance ongoing instrumental responding, a phenomenon termed Pavlovian–instrumental transfer (PIT). In its simplest form, PIT is assessed by presenting a CS for food while the subject is responding in extinction on a single lever for the same food. It has previously been shown that the nucleus accumbens, and particularly the core subdivision, is critical for this form of PIT (Hall *et al.*, 1999). However, behavioural studies have shown that PIT can be subdivided into a general, motivating effect of the Pavlovian CS, and a response-specific PIT effect, seen as a further enhancement when the Pavlovian and instrumental outcomes are the same (see Dickinson & Balleine, 1994). In the present study, rats received lesions to the core or shell of the nucleus accumbens before being tested on a response-specific PIT task. In the Pavlovian phase, one stimulus, CS(pel), was paired with pellet delivery, while a second stimulus, CS(suc), was paired with sucrose solution. The subjects were then trained to respond on two levers, with one lever producing pellets and the other producing sucrose. On test, lever-pressing was recorded in extinction while the stimuli were presented noncontingently. Control subjects ($n = 6$) showed a selective enhancement of lever-pressing on the lever paired with the same outcome as the Pavlovian CS; this is the response-specific PIT effect. Core-lesioned subjects ($n = 4$) showed a general enhancement of responding during the CS, but this was not specific to one response. Shell-lesioned animals ($n = 4$) showed no PIT. It is suggested that in this task, the shell is required for the ‘vigour’ and the core for the ‘direction’ of the potentiation of responding by a noncontingent, appetitive stimulus. This pattern closely resembles that previously observed for the potentiation of responding for conditioned reinforcement by psychostimulant drugs injected into the nucleus accumbens (Parkinson *et al.*, 1999b).

INTRODUCTION

Pavlovian CSs can have effects on operant responding when presented noncontingently, even when responding has never been associated with the CS. This was first demonstrated by Estes (1943; 1948), who found that noncontingent presentation of an appetitive CS, previously paired with food, would elevate the rate of instrumental responding for the same food in a test conducted in extinction — a phenomenon now known as Pavlovian–instrumental transfer (PIT). Estes used the simplest form of PIT, in which the instrumental outcome is the same as the Pavlovian US. As discussed in Chapter 1 (p. 26), PIT has since been subdivided behaviourally (see Dickinson, 1994, pp. 66–68; Dickinson & Balleine, 1994): the CS has a general, motivating effect (Dickinson & Dawson, 1987b; Balleine, 1994), but also potentiates, selectively, an action whose outcome is the Pavlovian US (Colwill & Rescorla, 1988; Colwill & Motzkin, 1994).

The nucleus accumbens (Acb) is an important neural site mediating the ability of Pavlovian CSs to invigorate and direct behaviour, as discussed in Chapter 1 (p. 46). It has previously been shown that PIT, in its simplest form, depends upon the nucleus accumbens core (AcbC), though not the shell (AcbSh)

(Hall *et al.*, 1999). However, it is not known how the AcbC and AcbSh contribute to response-specific PIT, or the potentiation of responses whose outcomes are unrelated to the US.

In the present experiment, the response-specific PIT effect was assessed in rats with lesions of the AcbC or AcbSh, using the experimental design of Colwill & Motzkin (1994, Experiment 2). To assess response-specific PIT, subjects were food-deprived, and two CSs were associated with different appetitive USs (chow pellets or sucrose solution). Next, two instrumental responses (left and right lever-presses) were trained for the two reinforcers used as USs in the Pavlovian phase. Finally, responding was tested in extinction while the CSs were presented noncontingently; response-specific PIT was inferred if one or both stimuli were capable of differentially affecting the two responses.

General PIT was assessed in the same subjects. In an attempt to detect general PIT in subjects who had already experienced several extinction sessions, the observation that PIT underlies the irrelevant incentive effect was used (Dickinson, 1986; Dickinson & Dawson, 1987b). Subjects were given further Pavlovian conditioning sessions, and retrained to respond on the lever producing the pellet outcome. They were shifted from a state of hunger to one of thirst, and their responding for pellets was again assessed in extinction while the CSs were presented. When subjects are thirsty, the CS for liquid sucrose solution predicts an outcome relevant to their current motivational state, and thus should produce strong Pavlovian conditioned motivation (Dickinson, 1986; Dickinson & Dawson, 1987b). General PIT was inferred if the CS for sucrose elevated responding for pellets. (Strictly, this assessment should also be made relative to an unpaired stimulus, though one was not available in the present experiment for technical reasons; in lieu of this, a comparison was made with responding in the interstimulus interval, but also with the CS for pellets, which was expected to suppress responding in thirsty animals; Balleine, 1994.) This test therefore relies on subjects' ability to use a CS to retrieve information about the US and assess its relevance to the current motivational state, as well as their capacity to show general (non-response-specific) PIT.

The experimental design is shown in Table 12.

Table 12. Design of the present experiment (S1, S2 = stimuli; L1, L2 = levers).

Training	Test 1 (Specific PIT)	Retraining	Test 2 (General PIT)
Hungry	Hungry	Hungry	Thirsty
S1 → pellet	S1: L1 v. L2	S1 → pellet	S1: L1
S2 → sucrose	S2: L1 v. L2	S2 → sucrose	S2: L1
L1 → pellet	ISI: L1 v. L2	L1 → pellet	ISI: L1
L2 → sucrose			
	Specific PIT inferred if one or both stimuli differentially affects the level of the two responses.		General PIT inferred if L1(S2) > L1(ISI).

METHODS

Six subjects (JP1, JP3, JP4, JP5, JP7, JP8) received excitotoxic lesions of the AcbC (see Chapter 2, p. 64, for coordinates) and had prior experience of an autoshaping task. Seven subjects (JP11, JP12, JP13, JP14, JP15, JP16, JP17) received lesions of the AcbSh and had prior experience of a simple visual discrimination task using a touchscreen. Six control subjects were used, of which two had received sham AcbC operations (JP6, JP9), two had sham AcbSh operations (JP19, JP20) and two were unoperated (JP2, JP10). No subjects had prior experience of the stimuli, reinforcers or responses used in the present task.

‘Response-specific’ and ‘general’ Pavlovian-instrumental transfer tests

Subjects were maintained at 85% of their free-feeding mass. Water was always available in the home cage during training except where stated. The two reinforcers used were 0.05 ml of 20% w/v (= 200 g/l = 0.58 M) sucrose so-

lution (the dipper was normally raised and was lowered briefly to collect liquid), and one 45-mg chow pellet (Rodent Diet Formula A/I, Noyes, Lancaster, NH). Stimulus S1 consisted of the left and right stimulus lights (2.8 W bulbs) above the levers, flashed together at 3 Hz. Stimulus S2 was a clicker relay operated at 10 Hz. A 2.8 W houselight was illuminated at all times.

Subjects were distributed evenly into the counterbalancing conditions listed in Table 13.

Table 13. Counterbalancing conditions for the response-specific PIT test.

Counterbalancing condition	Pavlovian	Instrumental	Sessions
0	S1→sucrose, S2→pellet	Left → sucrose, right → pellet	Begin with sucrose
1	S1→pellet, S2→sucrose	Left → sucrose, right → pellet	Begin with sucrose
2	S1→sucrose, S2→pellet	Left → pellet, right → sucrose	Begin with sucrose
3	S1→pellet, S2→sucrose	Left → pellet, right → sucrose	Begin with sucrose
4	S1→sucrose, S2→pellet	Left → sucrose, right → pellet	Begin with pellet
5	S1→pellet, S2→sucrose	Left → sucrose, right → pellet	Begin with pellet
6	S1→sucrose, S2→pellet	Left → pellet, right → sucrose	Begin with pellet
7	S1→pellet, S2→sucrose	Left → pellet, right → sucrose	Begin with pellet

Phase 1: Pavlovian training. Stimuli S1 and S2 were presented alternately in 2-min components, with no lever present. During each stimulus, the appropriate reinforcer was delivered on a RT 30-s schedule. Each component was separated from the next by a 2-min interstimulus interval (ISI). The session began with an ISI component and ended after 5 of each type of component had been presented. The stimulus–reinforcer assignment and the first reinforcer of the session were counterbalanced as shown in Table 13. Subjects were trained for 10 sessions with one session per day.

Phase 2: Instrumental training. For six 30-min sessions, animals were presented with a single lever that was reinforced on an RI schedule. No other stimuli were present. The lever used alternated across sessions, with half of the subjects receiving the pellet lever first and half the sucrose lever. The parameter of the RI schedule was 2 s for the first pair of sessions, 15 s for the second pair and 30 s for the third. For a further four sessions, both levers were present and reinforced on independent RI 30-s schedules. All sessions began with the insertion of the lever(s) and ended with lever retraction.

Phase 3: Pavlovian reminder. One further Pavlovian session was given, using the same schedule as Phase 1.

Phase 4: Instrumental extinction. As PIT is best observed after a degree of instrumental extinction has occurred (Dickinson *et al.*, 2000; A. Dickinson, personal communication, 7 May 1999), one 8-min session was given in which both levers were available but not reinforced.

Phase 5: Response-specific transfer test. Animals remained food-deprived for two sessions on the specific transfer test, in which both levers were available but not reinforced. Two-minute light and clicker stimuli were presented in alternation, with a 2-min ISI between each, until five of each stimulus had been presented. Assessing performance in the absence of the stimuli gave a measure of baseline lever-pressing. The session began with an ISI and lasted 40 min. As the stimulus presentation order was always ISI→S1→ISI→S2, half the rats received the stimulus associated with sucrose first, half received the stimulus indirectly associated with the left lever first, and half received the stimulus that occurred first in Pavlovian training; these three divisions were orthogonal (Table 14).

Phase 6: Retraining. Pavlovian retraining was given exactly as before for 3 sessions. This was followed by instrumental training in which only the pellet lever was present; three reinforced sessions were given using an RI 30-s schedule, followed by a single 5-min extinction session.

Phase 7: General transfer test. Once the animals had been fed after the final retraining session, they were placed on a 23-h water deprivation schedule with food freely available. On the next and subsequent day, they received a general transfer test in which only the pellet lever was available, though it was not reinforced. Three components were presented (light, clicker, no stimulus) in the same manner as for the specific transfer test, and responding was measured in each component.

As the specific and general PIT tests were conducted using the same subjects, it is important that the first test should not be able to bias the results of the second. This was the case: thus, retraining did not alter any of the Pav-

lovian or instrumental contingencies experienced by the subjects. The specific PIT test consisted of extinction trials to S1 and S2, but these were equal in number and duration. Furthermore, differential extinction on the two levers was not problematic, for the general PIT test was concerned with responding on a single lever. Finally, responding in extinction in the presence of S1 and S2 (in the specific PIT test) was unlikely to affect performance on the general PIT test: since no reinforcers were delivered, S1 and S2 could not become positive instrumental discriminative stimuli.

Table 14. Counterbalancing conditions for response-specific PIT test (continued from Table 13).

Counterbalancing condition	Begin with sucrose or pellet stimulus?	Begin with stimulus associated indirectly with left/right lever?	Begin with stimulus that occurred first or second in Pavlovian training?
0	sucrose	L	1
1	pellet	R	2
2	sucrose	R	1
3	pellet	L	2
4	sucrose	L	2
5	pellet	R	1
6	sucrose	R	2
7	pellet	L	1

RESULTS

Histology

Following histological analysis, the control group included six subjects (nos. JP2, JP6, JP9, JP10, JP19, JP20). Throughout the behavioural analyses, no differences were evident between subjects that had received sham AcbC surgery, sham AcbSh surgery, or no surgery; these subjects were therefore pooled to form a single sham group. In the core group, two rats were found to have lesions of the entire Acb (JP1, JP3) and were excluded from analysis, leaving four with bilateral core lesions only (JP4, JP5, JP7, JP8). In the shell group, one animal was found to have a septal lesion (JP14) and two to have no shell damage (JP16, JP17); these animals were excluded, leaving four with bilateral shell lesions (JP11, JP12, JP13, JP15).


Lesions of the AcbC encompassed most of the core subregion; neuronal loss and associated gliosis extended in an anteroposterior direction from approximately +2.5 mm to +0.5 mm relative to bregma, and did not extend ventrally or caudally into the ventral pallidum or olfactory tubercle. Damage to the ventromedial caudate–putamen was occasionally seen; damage to the AcbSh in these animals was restricted to the lateral edge of the dorsal shell or the superior edge of the lateral shell. Representative photomicrographs of AcbC lesions are shown in Figure 43 and Figure 44; schematics of the lesions are shown in Figure 45.

Lesions of the AcbSh encompassed the medial shell; neuronal loss and associated gliosis extended in an anteroposterior direction from approximately +2.2 mm to +1.0 mm relative to bregma. There was very little damage to the AcbC, the lateral septum, or the medial ventral pallidum. Representative photomicrographs of AcbSh lesions are shown in Figure 46 and Figure 47; schematics of the lesions are shown in Figure 48.


Unfortunately, the post-histological groups were not evenly counterbalanced. The core group were distributed evenly across counterbalancing conditions 1, 3, 5 and 7 (one rat per cell); thus they were counterbalanced for response/outcome assignment and stimulus presentation order, but not for stimulus/outcome assignment: all core-lesioned subjects received light→pellet and clicker→sucrose conditioning. The other groups were better counterbalanced. The shell group were in conditions 1, 4, 5 and 7. The sham group were in conditions 0, 2 (two rats), 4, 5 and 6.

Nucleus accumbens core photomicrographs (cresyl violet staining)


A. Sham-operated rat (low power)



B. Sham-operated rat (high power)



C. Core-lesioned rat (low power)



D. Core-lesioned rat (high power)






Figure 43. Lesions of the AcbC: photomicrographs of sections at approximately 1.2 mm anterior to bregma, stained with cresyl violet. **A & B:** sham-operated rat (ac, anterior commissure; CPu, caudate–putamen; AcbC, nucleus accumbens core; AcbSh, nucleus accumbens shell; LV, lateral ventricle). **C & D:** core-lesioned rat. Dotted lines show the extent of the lesion. There is tissue collapse within the lesion and the lateral ventricle is slightly expanded. **Left-hand panels** are low-magnification views (scale bars are 1 mm); **right-hand panels** are high-magnification views (scale bars are 0.1 mm). Arrowheads indicate the position of identical structures in corresponding pairs of photomicrographs.

Nucleus accumbens core photomicrographs (*NeuN immunocytochemical staining*)


A. Sham-operated rat (low power)



B. Sham-operated rat (high power)



C. Core-lesioned rat (low power)



D. Core-lesioned rat (high power)



Figure 44. Lesions of the AcbC: photomicrographs of sections at approximately 1.2 mm anterior to bregma, stained with NeuN antibody. **A & B:** sham-operated rat (ac, anterior commissure; CPu, caudate–putamen; AcbC, nucleus accumbens core; AcbSh, nucleus accumbens shell; LV, lateral ventricle). **C & D:** core-lesioned rat. Dotted lines show the extent of the lesion. **Left-hand panels** are low-magnification views (scale bars are 1 mm); **right-hand panels** are high-magnification views (scale bars are 0.1 mm). Arrowheads indicate the position of identical structures in corresponding pairs of photomicrographs.

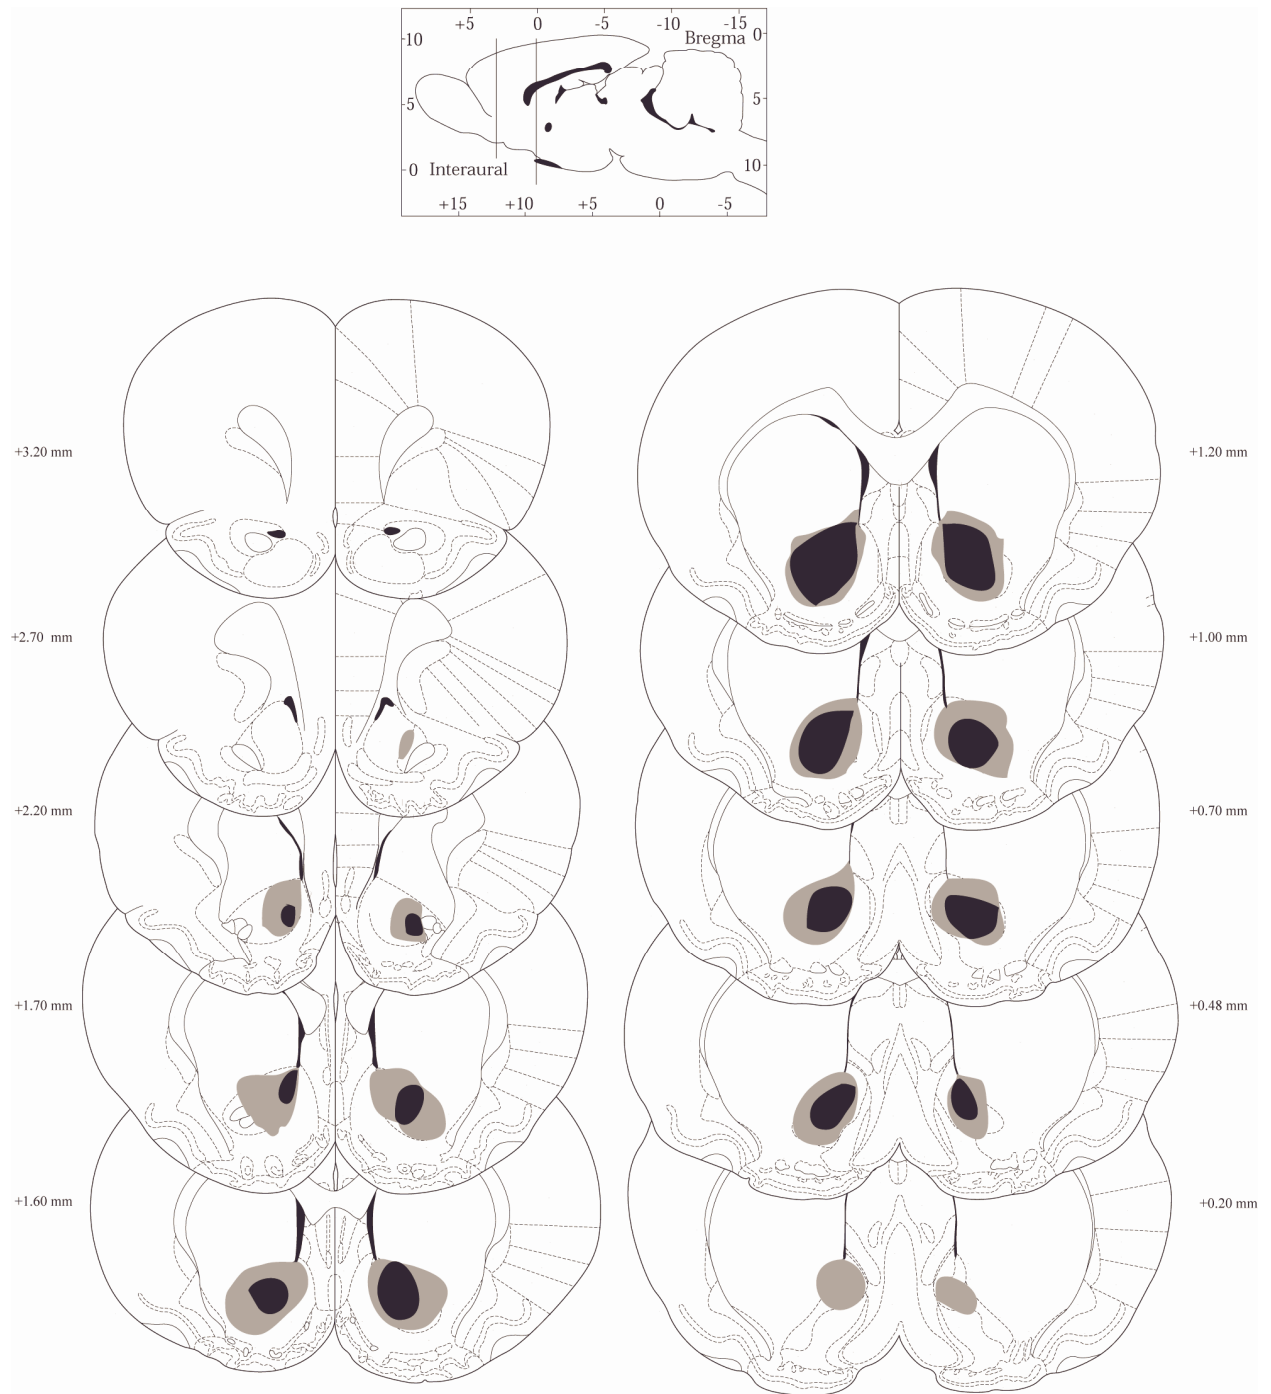

Nucleus accumbens core: schematic of lesions


Figure 45. Schematic of lesions of the AcCb (subjects JP4, JP5, JP7, JP8). Grey shading indicates the extent of the largest area of neuronal loss, and black the smallest. Diagrams are taken from Paxinos & Watson (1998).

Nucleus accumbens shell photomicrographs (cresyl violet staining)


A. Sham-operated rat (low power)



B. Sham-operated rat (high power)



C. Shell-lesioned rat (low power)



D. Shell-lesioned rat (high power)






Figure 46. Lesions of the AcbSh: photomicrographs of sections at approximately 1.0 mm anterior to bregma, stained with cresyl violet. **A & B:** sham-operated rat (ac, anterior commissure; CPu, caudate–putamen; AcbC, nucleus accumbens core; AcbSh, nucleus accumbens shell; LV, lateral ventricle). **C & D:** shell-lesioned rat. Dotted lines show the extent of the lesion; the lesioned area has collapsed and there is some ventricular expansion. **Left-hand panels** are low-magnification views (scale bars are 1 mm); **right-hand panels** are high-magnification views (scale bars are 0.1 mm). Arrowheads indicate the position of identical structures in corresponding pairs of photomicrographs.

Nucleus accumbens shell (NeuN immunocytochemical staining)

A. Sham-operated rat (low power)



B. Sham-operated rat (high power)



C. Shell-lesioned rat (low power)



D. Shell-lesioned rat (high power)




Figure 47. Lesions of the AcbSh: photomicrographs of sections at approximately 1.0 mm anterior to bregma, stained with NeuN antibody. **A & B:** sham-operated rat (ac, anterior commissure; CPu, caudate–putamen; AcbC, nucleus accumbens core; AcbSh, nucleus accumbens shell; LV, lateral ventricle). **C & D:** shell-lesioned rat. Dotted lines show the extent of the lesion. **Left-hand panels** are low-magnification views (scale bars are 1 mm); **right-hand panels** are high-magnification views (scale bars are 0.1 mm). Arrowheads indicate the position of identical structures in corresponding pairs of photomicrographs.

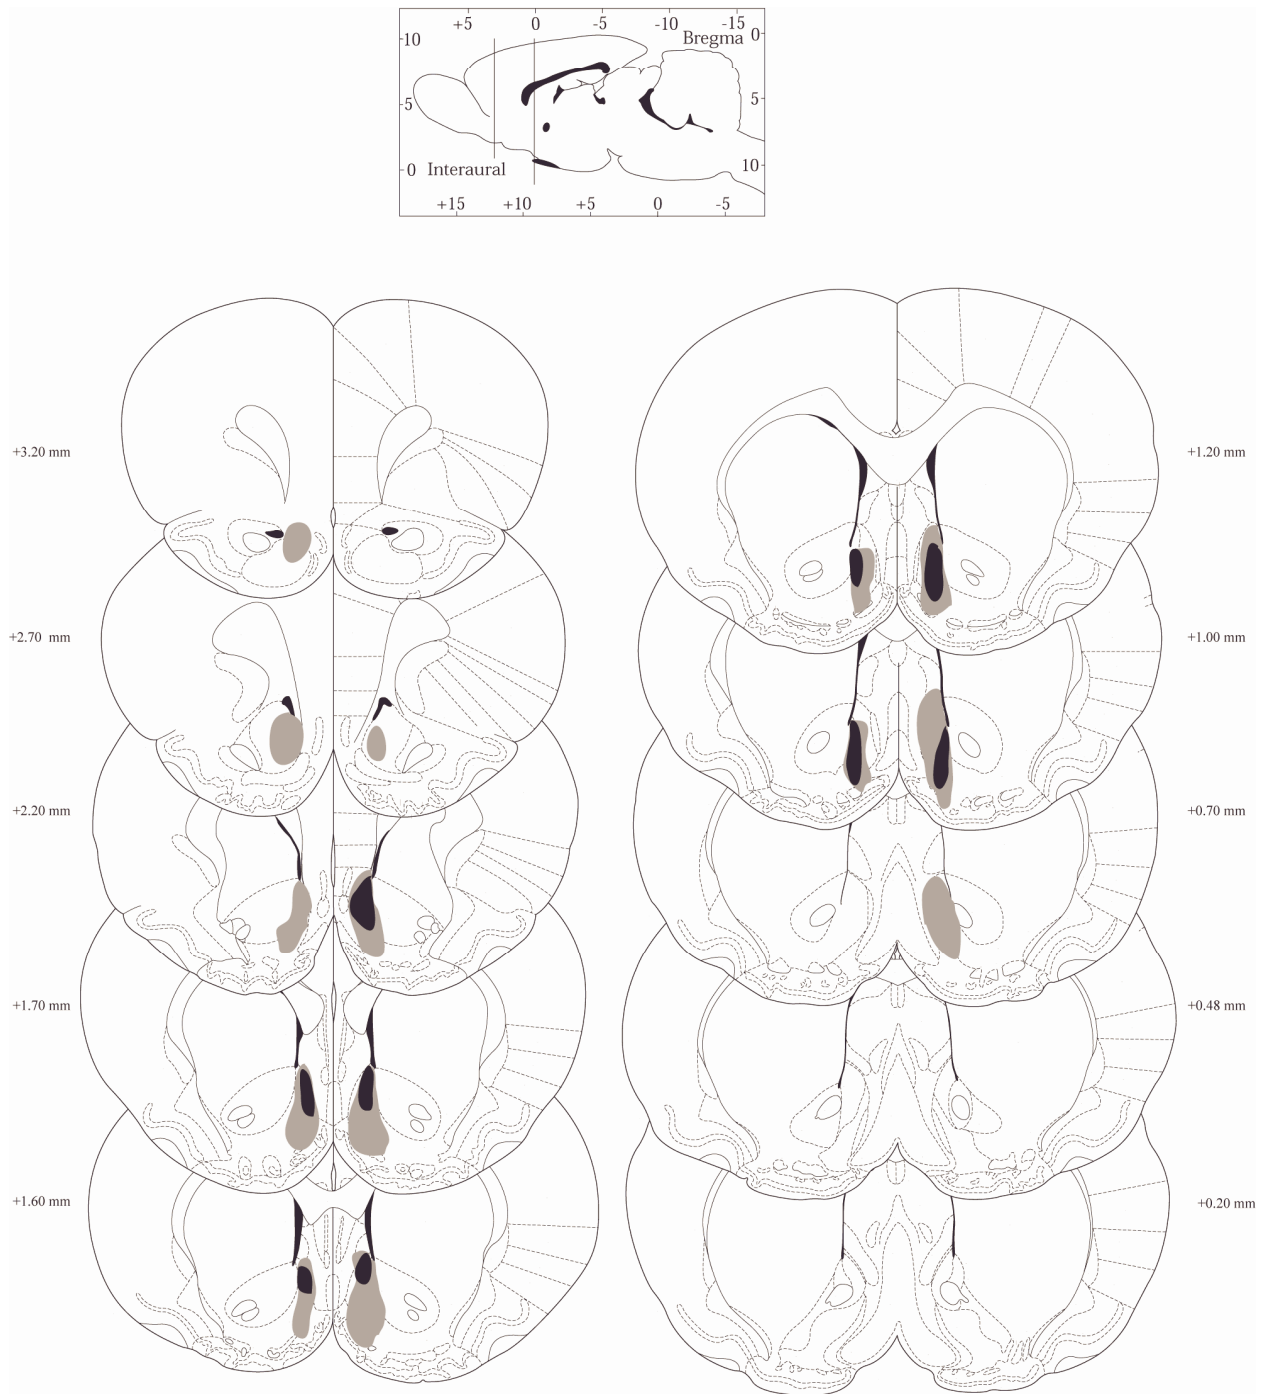
Nucleus accumbens shell: schematic of lesions

Figure 48. Schematic of lesions of the AcbSh (subjects JP11, JP12, JP13, JP15). Grey shading indicates the extent of the largest area of neuronal loss, and black the smallest. Diagrams are taken from Paxinos & Watson (1998).

Pavlovian training

The groups did not differ in their approach behaviour during Pavlovian training. Figure 49 shows approach to the food alcove during the two stimuli, CS(pel) and CS(suc), relative to the ISI. In this task, the measure of approach behaviour is not a pure measure of conditioning (be it Pavlovian or instrumental approach): as food is delivered during the stimuli, approach may reflect unconditioned responding.

Analysis of ratios of responding from initial training sessions using the model $\text{group}_3 \times (\text{stimulus}_2 \times \text{session}_{10} \times S)$ demonstrated a main effect of session ($F_{9,99} = 3.328, p = .001$) and stimulus ($F_{1,11} = 19.506, p = .001$), with nosepoking being greater during the sucrose CS. There was no session \times stimulus interaction ($F_{9,99} = 1.745, \text{NS}$). However, there was no significant effect of group, and no interactions involving group ($F_s < 1.22$).

Nor did the groups differ during the ‘reminder’ session, for which a separate ANOVA was conducted. Again, responding was higher during the sucrose stimulus ($F_{1,11} = 8.553, p = .014$) but there were no group differences (group: $F_{2,11} = 1.571, \text{NS}$; group \times stimulus: $F_{2,11} = 2.14, \text{NS}$).

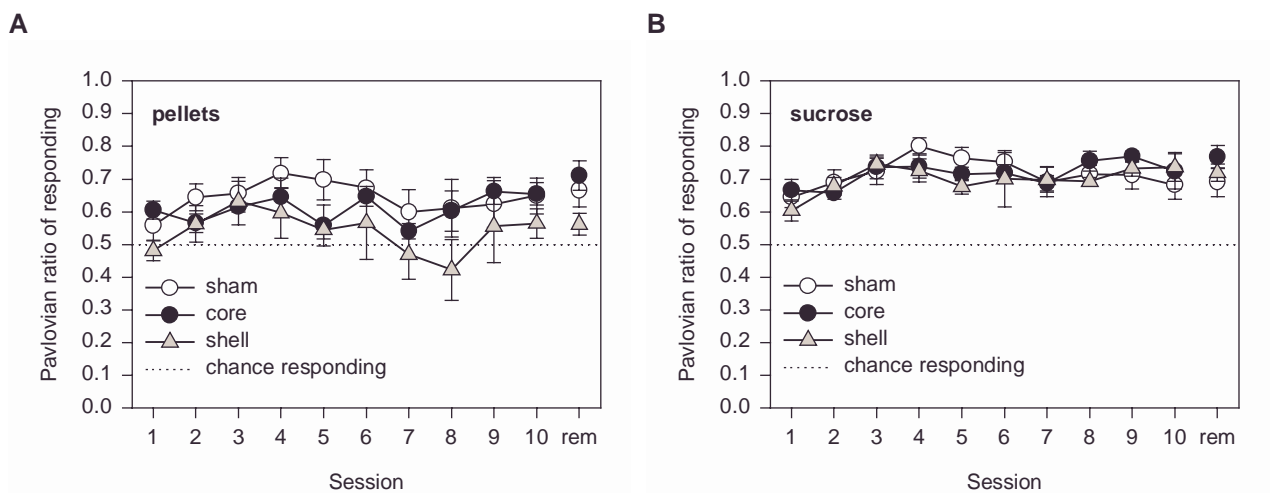


Figure 49. Approach to the food alcove during Pavlovian training, measured by the nosepoke detector. The ratio of responding is the approach time during reinforcer delivery, as a proportion of approach during reinforcer delivery plus approach during the ISI. It is not a pure measure of conditioning, as unconditioned responding to the food may contribute. **A:** Approach during pellet delivery. **B:** Approach during sucrose delivery. The reminder session is indicated by *rem*. The two graphs show approach relative to the same ISI data.

Instrumental training

All groups acquired the lever-press responses at the same rate (Figure 50A). The total number of lever-presses in each session was square-root transformed and data from training sessions 1–10 were analysed using the model $\text{group}_3 \times (\text{session}_{10} \times S)$. This showed a main effect of session ($F_{7,133,70.785} = 42.528, \tilde{\epsilon} = .715, p < .001$) but no group differences ($F_s < 1.23, \text{NS}$).

The core group displayed a slightly stronger preference for the pellet lever than the other two groups. From session 7, when two levers were available, the preference for the sucrose lever was calculated as (sucrose responses) \div (total responses) and subjected to ANOVA. These data are shown in Figure 50B, where it can be seen that all groups responded almost equally on both levers by the end of training (proportions close to 0.5). However, the ANOVA demonstrated a main effect of group ($F_{2,11} = 4.296, p = .042$), with no effect of session and no interaction ($F_s < 1$). Pairwise comparisons with a Sidak correction suggested that this difference was due to the core group having *lower* preference scores than shams (i.e.

preferring the pellet lever more), $p = .052$, with no difference between the shell group and shams ($p = .216$) or between the core and shell groups ($p = .841$).

The rate of responding in extinction did not differ between the groups (Figure 50A), though the core group continued to prefer the pellet lever (Figure 50B). Separate one-way ANOVAs were conducted for the extinction session, which demonstrated no difference in total lever-pressing between the groups ($F_{2,11} = 1.171$, NS). However, the preference for the pellet lever in the core group increased: there was a significant main effect of group ($F_{2,11} = 4.141$, $p = .046$), and Dunnett's test showed that the core group had lower preference scores than shams ($p = .03$) but the shell group did not ($p = .248$).

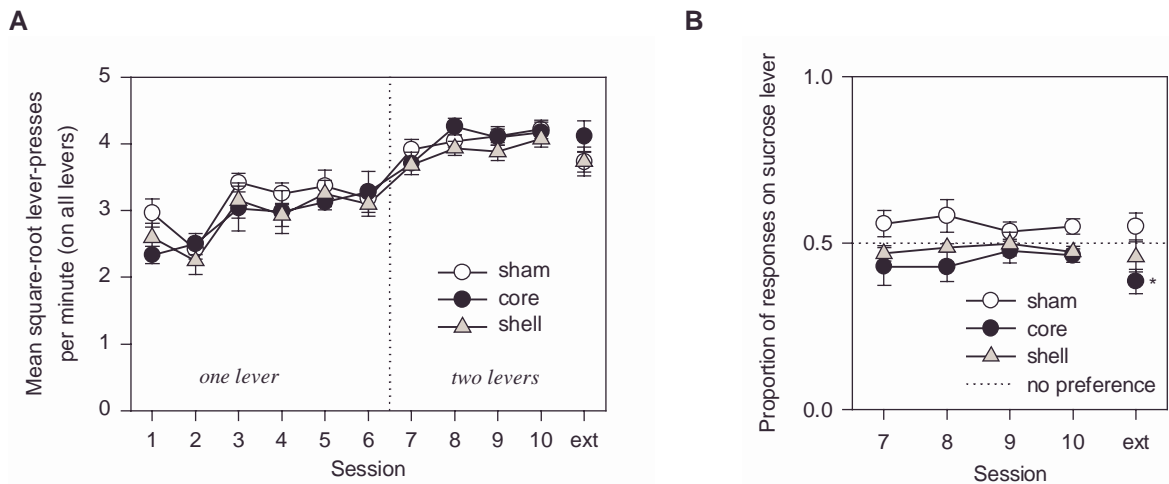


Figure 50. Acquisition of instrumental responding. **A:** Responses per minute, summed over all levers present. The session length was 30 min; the schedule progressed from RI 2 s to RI 30 s as described in the Methods. From session 7 onwards, two levers were concurrently available; *ext* indicates the extinction session. **B:** Proportion of responses made on the sucrose lever (* $p < .05$ relative to shams on the extinction day).

Response-specific PIT

Response rates for the two levers during each stimulus condition are shown in Figure 51, with rates of nose-poking in Figure 52.

Lever-pressing

The sham group displayed a response-specific PIT effect, in that the CS for pellets selectively enhanced responding on the pellet lever. The core group, which preferred the pellet lever slightly, displayed PIT, but this was not specific: the CS for pellets potentiated responding on both levers. The shell group displayed no PIT. The CS for sucrose was less effective than the CS for pellets in producing PIT, across the groups.

Lever-press data were square-root transformed before analysis. As the groups were not evenly counterbalanced for stimulus/outcome assignment, a four-way ANOVA including this term was first performed. This failed to demonstrate any effect of the stimulus/outcome assignment ($F_s < 1.477$, $p > .255$); consequently this term was removed from further analyses.

An ANOVA using the model $\sqrt{(\text{lever-presses})} = \text{group}_3 \times (\text{response}_2 \times \text{stimulus}_3 \times S)$ revealed a response \times stimulus \times group interaction ($F_{4,22} = 3.741$, $p = .018$), in addition to a main effect of response ($F_{1,11} = 5.15$, $p = .044$) and stimulus ($F_{2,22} = 7.646$, $p = .003$). No other terms were significant ($p > .1$). The response \times stimulus \times group interaction was analysed further by considering simple interaction terms; that is, testing for a response \times stimulus interaction in each group.

In the sham group there was a significant response \times stimulus interaction ($F_{2,10} = 8.312, p = .007$), indicating that the pattern of responding across the two levers was differentially affected by the stimuli and implying a response-specific PIT effect. Further analysis showed that the stimulus \times response interaction was due to the CS(pel) selectively potentiating responding on the pellet lever (simple effect of stimulus for the pellet lever, $p < .001$, but not for the sucrose lever, $p = .071$; pairwise comparisons for the pellet lever showed that only the CS(pel) elevated responding, $p = .005$); neither stimulus affected responding on the sucrose lever ($F_{2,10} = 3.489, p = .071$).

In the core group, there was no response \times stimulus interaction ($F < 1$), though these animals demonstrated both a preference for the pellet lever (main effect of response, $F_{1,3} = 26.05, p = .015$) and some response-independent PIT (main effect of stimulus, $F_{2,6} = 24.273, p = .001$; indeed, the pellet CS elevated responding on the sucrose lever, $p = .018$). Pairwise comparisons showed that responding during the CS(pel) was significantly higher than during the ISI ($p = .002$). It was also higher than during CS(suc), though this did not reach significance ($p = .08$); there was no difference in responding between the ISI and CS(suc) ($p > .5$).

The shell group demonstrated no response \times stimulus interaction ($F_{2,6} = 3.159, p = .116$), nor a main effect of either response or stimulus ($F_s < 1$). The lack of an effect of the CSs was not due to differences in the baseline level of responding that obscured PIT; comparison of ISI responding between shell- and sham-lesioned rats using the model group \times (response \times S) revealed no differences ($F_s \leq 1.214, NS$).

These results therefore demonstrate a Pavlovian–instrumental transfer effect in normal rats that is response-specific in that Pavlovian CSs for two reinforcers differentially affected lever-pressing for those reinforcers. Shell-lesioned animals demonstrated no Pavlovian–instrumental transfer, while core-lesioned animals demonstrated transfer, but this transfer lacked response specificity.




Figure 51. Lever-pressing during the specific PIT test. Sham-operated controls exhibited a specific Pavlovian–instrumental transfer effect, with the Pavlovian CS for pellets selectively potentiating responding on the lever trained with the pellet outcome. Core-lesioned animals exhibited a PIT effect in response to the CS for pellets, but the potentiation was not response-specific. Shell-lesioned animals exhibited no transfer effect. (** $p < .01$; *** $p < .001$.)

Supplemental analysis

For the purpose of comparison with a study of PIT by Hall *et al.* (1999), who used a single test session and a design with a single lever that produced a pellet reinforcer, data from only the first session of response-specific PIT testing were taken, and responding on the pellet lever was analysed in isolation, using the model $\text{group}_3 \times (\text{stimulus}_3 \times S)$. This revealed a main effect of stimulus ($F_{1,777,19,548} = 4.929$, $\tilde{\epsilon} = .889$, $p = .021$); the group \times stimulus interaction escaped significance ($F_{3,554,19,548} = 2.542$, $\tilde{\epsilon} = .889$, $p = .078$) and the main effect of group was not significant ($F < 1$, NS). Overall, the pellet CS elevated responding relative to the ISI ($p = .026$) but the sucrose CS had no effect ($p = .351$). This was also true of the sham group considered alone (pellet CS *v.* ISI, $p = .004$; sucrose CS, $p = .85$). However, in this analysis, no effect of either stimulus was detectable for the core or the shell groups ($F_s < 1$, NS).

Nosepoking

Neither CS affected the rate of nosepoking in core- and shell-lesioned subjects. In the sham group, there was a tendency for the CS(pel) to elevate nosepoking, but this was ambiguous statistically. To analyse nosepoking, the rate of nosepoking was calculated and subjected to a square-root transform to improve homogeneity of variance before an ANOVA was performed using the model $\text{group}_3 \times (\text{stimulus}_3 \times S)$. This revealed a significant stimulus \times group interaction ($F_{4,22} = 4.081$, $p = .013$). Simple effects analysis showed that nosepoking differed among the three stimulus conditions in the sham group ($F_{2,10} = 4.304$, $p = .045$); though no condition was different from any other by *post hoc* pairwise comparisons ($p > .18$), inspection of Figure 52 suggests that the effect was due to elevation of nosepoking by the CS for pellets. In the core and shell groups, there was no effect of the stimulus (core: $F_{2,6} = 3.82$, $p = .085$; shell: $F_{2,6} = 2.334$, NS).




Figure 52. Nosepoking during the specific PIT test. While the stimulus condition affected the rate of nosepoking in the sham group, as detected by a main effect in an ANOVA, no one stimulus condition was significantly different from another by *post hoc* pairwise comparisons. The stimuli did not affect the rate of nosepoking in the core- or shell-lesioned rats.

Retraining

Retraining data are shown in Figure 53.

The groups did not differ in their approach behaviour during Pavlovian retraining. Analysis of the ratios of responding for the Pavlovian sessions was performed using the model $\text{group} \times (\text{session}_3 \times S)$. Once more, this showed greater approach while the CS(suc) was on and sucrose was being presented than dur-

ing the pellet stimulus ($F_{1,11} = 7.394, p = .02$), but no other terms were significant (closest to significance was session \times group, $F_{4,22} = 2.229, p = .099$).

The three groups reacquired the instrumental response at approximately the same rate, but the core-lesioned rats maintained a higher rate of responding in the subsequent extinction session. Analysis of square-root-transformed lever-press data for the reinforced training sessions showed that responding increased over the three sessions ($F_{1,288,14,173}, \tilde{\epsilon} = .644, p = .02$), but no differences between groups were significant (group: $F_{2,11} = 2.356, p = .141$; group \times session: $F_{2,577,14,173} = 1.173, \tilde{\epsilon} = .644, \text{NS}$). However, the core-lesioned group responded significantly more than shams on the extinction day (univariate ANOVA, $F_{2,11} = 4.613, p = .035$; Dunnett's test showed that the core group responded more than shams, $p = .024$, but the shell group did not, $p = .767$).




Figure 53. Retaining data before the general PIT test. Panels **A & B** show ratios of responding for the two stimuli in the Pavlovian retraining sessions; as for Figure 49, the ratio of responding is calculated from the time spent approaching the food alone using the formula $CS/(CS+ISI)$. Panel **C** shows responding on the pellet lever (the only lever available) during instrumental retraining and a further extinction session, plotted to the same scale as Figure 50A (* $p < .05$ relative to shams).

General PIT

Lever-pressing was not affected by stimulus presentation in the general transfer test, and the core group responded more than the other groups (Figure 54). Following square-root transformation, an ANOVA was performed using the model $group_3 \times (stimulus_3 \times S)$. This showed that the effect of stimulus presentation did not reach significance ($F_{1,54,16,94} = 3.366, \tilde{\epsilon} = .77, p = .069$) and that there was no stimulus \times group interaction ($F < 1$), though there was a main effect of group ($F_{2,11} = 5.861, p = .019$). The group difference was due to the core group responding more than the other two ($p = .017$), which did not differ from each other ($p > .24$).

However, the Pavlovian stimuli did affect the rate of nosepoking (Figure 55): the CS for sucrose elevated nosepoking relative to the ISI, while the CS for pellets was less effective. Analysis of square-root-transformed nosepoke rates revealed a main effect of stimulus ($F_{2,22} = 8.766, p = .002$), though no main effect of group ($F < 1$) and no stimulus \times group interaction ($F_{4,22} = 2.072, p = .119$). The effect of the stimulus condition could be attributed to greater responding during the sucrose stimulus than the ISI ($p = .007$), with responding during the pellet stimulus at an intermediate level (pellet stimulus *v.* ISI, $p = .058$; pellet *v.* sucrose stimulus, $p = .211$; overall means, in units of square-root responses per minute: ISI 1.393 ± 0.123 , pellet stimulus 1.623 ± 0.174 , sucrose stimulus 1.924 ± 0.11).




Figure 54. Lever-pressing during the first general transfer test. Responding was not affected by the stimuli, though the core group responded more than the other two groups.




Figure 55. Nose-poking during the general transfer test. The CS(suc) increased the rate of nose-poking over the ISI, but there were no differences between the groups.

On the supposition that transfer to nose-poking rather than lever-pressing occurred because nose-poking was the current prepotent response, further training was given to encourage lever-pressing on test. The rats were returned to the food-deprivation state and given a single Pavlovian training session followed by 8 instrumental sessions with the pellet lever (RI-30s schedule), using the methods described previously. After this, the two-day general transfer test was repeated with subjects water-deprived.

However, this second test was not successful. By this stage the subjects had had extensive experience with extinction sessions (both of the Pavlovian stimuli and the instrumental responses), and responded at very low levels. Responding increased across the instrumental retraining sessions, and inspection of the data suggested that the core group maintained a higher rate of responding, but this was not a significant difference ($p \geq .147$). During the second general transfer test, core-lesioned animals responded more than shams (effect of group: $F_{2,11} = 9.246$, $p = .004$; pairwise comparisons established that the core group differed from the other two groups, which did not differ from each other). However, there was no effect of stimulus presentation on lever-pressing ($F_s \leq 1.571$, NS) or the rate of nose-poking ($F_s < 1.23$, NS).

DISCUSSION

Despite the small number of subjects in the present experiment, response-specific PIT was successfully demonstrated in the sham group: the pellet CS selectively enhanced responding on the lever producing pellets. In core-lesioned subjects, PIT was observed but was not response-specific, while shell-lesioned subjects exhibited no PIT.

‘General’ PIT, the potentiation of responding for one reinforcer by a CS for another, was not successfully obtained in the sham group. As this effect has reliably been observed in other, broadly similar, experimental designs (Dickinson, 1986; Dickinson & Dawson, 1987a; Dickinson & Dawson, 1987b; Balneine, 1994), two features of the present design probably contributed to this failure. Firstly, the results of the response-specific PIT test suggested that the CS for pellets was more effective than the CS for sucrose (the CS for sucrose had no effect on any subjects in this test). Despite the attempt to make the CS for sucrose more salient by rendering the subjects thirsty, this CS was ineffective during the general PIT test except to elevate nose-poking slightly. Had the CS for pellets been presented to hungry subjects responding for sucrose, an effect might have been observed. Indeed, general PIT *was* observed in core-lesioned subjects during the response-specific transfer test, in that the CS for pellets elevated responding on the sucrose lever. Secondly, the general PIT test was conducted after subjects had experienced several extinction sessions (both for the CSs and the responses) as part of the response-specific PIT test. The attempt to conserve subjects was perhaps overly ambitious, and general PIT may be easier to demonstrate in experimentally naïve subjects. This is *not* to imply that the effect is biologically unimportant (as discussed in Chapter 1, p. 27, it underlies the irrelevant incentive effect, probably of great functional significance), but simply that the extinction procedure used to demonstrate the effect guarantees that it will be ephemeral (cf. conditioned reinforcement; Mackintosh, 1974, p. 237).

The psychological basis of response-specific PIT

The present results support some, but not all, previous theories of the psychological basis of response-specific PIT. Several slightly different experimental designs have been used to demonstrate this effect (reviewed by Colwill & Motzkin, 1994). For example, Colwill & Rescorla (1988) trained two groups of subjects, each experiencing a single CS — for one group, the CS was paired with pellets, and for the other group, it was paired with sucrose (see Chapter 1, p. 26). The subjects were then trained to press a lever for pellets and pull a chain for sucrose in separate sessions (of course, the experiment was counterbalanced in this respect). As discussed by Dickinson (1994, p. 67), this meant that subjects learned to press the lever at a time when the contextual cues were associated with pellets and pull the chain when these cues were paired with sucrose solution. Consequently, the presentation of the Pavlovian CS on test may have helped to reinstate the conditions under which one of the actions was trained. (This explanation emphasizes the role of the stimuli that *elicit* instrumental responses, not the consequences of those responses.)

The present design was essentially that of Colwill & Motzkin (1994, Experiment 2), and the results support their conclusions regarding the psychological basis of the effect. As a within-subjects design was used, all animals experienced two Pavlovian CSs paired with two different reinforcers. These CSs were trained in alternation. Furthermore, the two instrumental responses were trained concurrently. This design minimizes differential contextual associations of the Pavlovian CSs, the instrumental responses, and the reinforcers. It is therefore less obvious that the CS reinstated the conditions under which the action was trained. The alternative, more likely explanation of the present data is that the CS potentiates actions based on a comparison of the US with the *outcome* of the instrumental response, as argued by Colwill & Motzkin (1994).

One other feature of the present behavioural results is worth noting. While some previous studies have found that a Pavlovian CS exerts its response-specific effect by *depressing* responses that do *not* share an outcome with the CS (e.g. Colwill & Rescorla, 1988; Colwill & Motzkin, 1994), the present results provide further evidence that a CS can selectively potentiate responses with which it does share an outcome (see also Baxter & Zamble, 1982). It is not at present clear why these difference was found, particularly as the present experiment was very similar in design to that of Colwill & Motzkin (1994), and suggests that CS/S^D differences may not be as critical as previously suggested in determining the direction of the effect (cf. Colwill & Rescorla, 1988). Possible explanations include differences in the rate of baseline responding, the degree of food deprivation, and the degree to which transfer occurs to behaviours other than instrumental responding, but this remains an area for further investigation.

The contribution of the Acb to PIT

It is difficult to draw a clear picture of the role of the Acb in PIT from the experiments conducted to date, as some studies appear contradictory.

The present experiment suggests that the AcbSh is required for PIT *per se*, perhaps providing the ‘vigour’ of PIT, while the AcbC is required to ‘direct’ this potentiation to a particular response when that response shares an outcome with the CS. These results provide further support for the claim that the Acb is critically involved in the impact of Pavlovian CSs upon behaviour (see Chapter 1 and Parkinson *et al.*, 2000a). It seems unlikely that these deficits were due to a failure to discriminate the two instrumental responses, as both core-lesioned (e.g. Chapter 7; Parkinson *et al.*, 1999b) and shell-lesioned (e.g. Parkinson *et al.*, 1999b) rats have been shown able to discriminate two levers for the purposes of responding. The pre-existing preference of the core group for the pellet lever undoubtedly complicates interpretation a little, but cannot easily explain the lack of response specificity in this group; the preference did not lead to a ‘ceiling effect’, for at response rates of ~4/min, core-lesioned subjects were certainly not responding maximally on the pellet lever (they responded at rates of ~16/min during instrumental acquisition, for example).

These results closely resemble the effects of core/shell lesions on the potentiation of responding for conditioned reinforcement by intra-Acb amphetamine (temporarily designated ‘amphetamine potentiation of conditioned reinforcement’, APCR). Parkinson *et al.* (1999b) showed that shell lesions abolished APCR, while core lesions removed the response selectivity of APCR without abolishing APCR itself.

The present results also show some correspondence to those of Corbit & Balleine (2000a), who found that AcbSh lesions abolished transfer in a variant of the response-specific PIT procedure. They found no effect of AcbC lesions on PIT, although only a single lever was present at any one time during their test, which may therefore have been less sensitive to deficits in response specificity (or in the ability to switch between responses as a result of CS presentation). An additional procedural difference was that Corbit & Balleine used a ratio schedule to demonstrate PIT. As discussed in Chapter 1 (p. 28), Lovibond (1983) showed that simple PIT may have a different psychological basis under ratio and interval schedules (possibly relating to the relative contributions of habitual and goal-directed behaviour; see Chapter 1), and it is not yet known how this relates to the involvement of the AcbC.

It is not so easy to reconcile the present data with those of Hall *et al.* (1999), who tested rats with a ‘simple’ PIT task, testing elevation of responding on a single lever by a CS for the same outcome. Hall *et al.* found that shell lesions had no effect on PIT, while core lesions completely abolished the effect. Intuitively, response-specific PIT has much in common with simple PIT: the response-specific test is the simple PIT test with another response available. On the basis of the results of Hall *et al.*, it would be expected

that AcbC lesions would abolish PIT entirely. If both sets of results are accepted, the puzzling conclusion is that response-specific PIT engages a (core-independent) process that does not contribute significantly to simple PIT (because core-lesioned animals showed some PIT in the former situation, but not the latter), but that this extra process is not response specificity itself (as the core-lesioned subjects did not show response specificity). A similar argument may be made from the finding that shell lesions impaired PIT in the present study, but not in that of Hall *et al.* (1999).

However, procedural differences do exist between the two studies. Hall *et al.* (1999) used only a pellet lever and a CS for pellets, and tested over a single 30-min session. In an attempt to see if this difference accounted for the discrepant findings, data from the present study were analysed in an analogous manner (p. 141). This analysis, while detecting PIT in the sham group, failed to detect PIT in the core or shell groups. It is possible, therefore, that the two-day test is more sensitive, and that this accounts for the detection of a PIT effect in the core group, though this cannot explain the differences in findings for the shell group. Perhaps measuring an additional response, as in the present experiment, simply increases the power to detect PIT. Indeed, as Figure 51 (p. 140) shows, the greatest PIT effect observed in the core-lesioned subjects was elevation of responding on the sucrose lever by the pellet CS! In the study of Hall *et al.* (1999), the CS did elevate the rate of one other behaviour, nose-poking in the food alcove — though even using this measure, core-lesioned subjects were impaired relative to shams (J. Hall, personal communication, 8 June 1999).

Additionally, technical failings of the present study must be taken into account. This experiment was based on a small number of subjects (sham 6, core 4, shell 4); though this does not alter any of the conclusions regarding these subjects, it brings a sense of caution to the interpretation of the results as representative of all sham-, core-, or shell-lesioned rats. Also, following histological analysis, the counterbalancing of the groups was incomplete. While an attempt was made to detect bias resulting from this failure of counterbalancing (p. 139), and none was found, failure to find any effects of the counterbalancing conditions may simply have been due to low statistical power and the ‘unbalanced’ counterbalancing may have contributed in some way to the results.

To summarize, while the present results are consistent with work concerning the role of the AcbC and AcbSh in APCR, surprising differences from previous studies of simple PIT emerged. As these differences suggest that PIT operates in a highly counter-intuitive manner, it would be well worth while replicating the present study with larger group sizes to give more effective counterbalancing, perhaps with a larger sucrose reinforcer in order to observe an effect of the CS for sucrose.

The relationship between PIT and conditioned reinforcement

Neither the AcbC nor the AcbSh appear to contribute to the basic phenomenon of conditioned reinforcement; however, they are both critically involved in the artificial phenomenon of APCR (Parkinson *et al.*, 1999b). Wyvell & Berridge (2000) have found that intra-Acb amphetamine potentiates PIT, implying that intra-Acb amphetamine has effects that cannot be explained solely in terms of conditioned reinforcement. It may be fruitful to ask whether the converse is true: can the contribution of the Acb to APCR be explained in terms of PIT, or do both phenomena need to be subsumed within a wider description?

PIT is clearly not analogous to conditioned reinforcement itself. As discussed in Chapter 1 (p. 27), general PIT does not affect choice behaviour (unlike CRf), although once a CS has been earned in a conditioned reinforcement task, it might be capable of boosting responding through PIT. Response-specific PIT might contribute to CRf (though this would require more than first-order associations; see Chapter 1, p. 31), but there is no direct evidence for this suggestion. Furthermore, PIT and CRf have been dissoci-

ated neurally; lesions of the CeA impair simple PIT (Hall *et al.*, 1999), but do not impair CRf (though the effect of intra-accumbens amphetamine upon CRf is abolished; Robledo *et al.*, 1996). Similar results have been reported by Killcross *et al.* (1998), using a task in which prolonged presentation of a putative conditioned reinforcer did indeed produce an elevation of responding (interpretable as PIT); this elevation was sensitive to CeA lesions but the CRf effect itself was not. Conversely, lesions of the BLA, which impair CRf (Cador *et al.*, 1989; Burns *et al.*, 1993), do not affect simple PIT (Killcross *et al.*, 1998; Hall *et al.*, 1999). Finally, lesions of the AcbC or AcbSh do not impair the acquisition of a new response with CRf (Parkinson *et al.*, 1999b), but these regions contribute to PIT (in a way that is still not completely clear: present experiments; Hall *et al.*, 1999; Cardinal *et al.*, 2000a; Corbit *et al.*, submitted).

However, there is a striking match between the neural bases of PIT and APCR. The present results suggest that the AcbSh is required for PIT *per se*, while the AcbC is not required for PIT but is required to ‘direct’ this potentiation to a particular response. Similarly, Parkinson *et al.* (1999b) showed that shell lesions abolished APCR, while core lesions removed only the response selectivity of APCR. The analogy may be continued: APCR depends upon Acb dopamine (Taylor & Robbins, 1986; Cador *et al.*, 1991; Wolterink *et al.*, 1993), while noncontingent presentation of an appetitive CS elevates Acb dopamine (specifically in the AcbC; Bassareo & Di Chiara, 1999; Ito *et al.*, 2000). PIT may also involve Acb dopamine, as it is abolished by systemic dopamine antagonists (Dickinson *et al.*, 2000) and enhanced by intra-Acb amphetamine (Wyvell & Berridge, 2000). Both APCR (Robledo *et al.*, 1996) and PIT (Hall *et al.*, 1999) depend on the CeA, probably because the CeA influences Acb dopamine via the VTA (see Chapter 1, pp. 43/47/49). Furthermore, lesions of the BLA remove the source of information to the Acb regarding conditioned reinforcement that determines which lever APCR acts upon (Cador *et al.*, 1989; Burns *et al.*, 1993); similarly, BLA lesions impair the response selectivity of PIT (Blundell & Killcross, 2000a) but do not abolish the basic PIT effect (Hall *et al.*, 1999; Blundell & Killcross, 2000a). Core lesions can sometimes abolish PIT (Hall *et al.*, 1999), and they also abolish APCR, in that the ability of amphetamine to potentiate responding for a conditioned reinforcer in a selective manner is lost, though amphetamine still potentiates responding in a nonselective manner in AcbC-lesioned animals (Parkinson *et al.*, 1999b). Shell lesions abolish APCR (Parkinson *et al.*, 1999b) and can abolish PIT (present experiments; Corbit & Balleine, 2000a), though not in all tasks (Hall *et al.*, 1999). Thus, though ambiguities remain, it may be reasonable to suppose that APCR reflects artificial activation of the system by which *noncontingent* Pavlovian CSs normally increase the probability of instrumental responses (PIT). This system appears to play a minor role in responding for CRf under normal situations (thus, responding for conditioned reinforcement survives AcbC and AcbSh lesions; Parkinson *et al.*, 1999b), possibly reflecting the fact that typical CRf experiments use brief conditioned reinforcers that cannot significantly potentiate responding via PIT.

Finally, as the Acb is also necessary for autoshaping (the AcbC, but not the AcbSh; Parkinson *et al.*, 2000c), a task in which the response is Pavlovian locomotor approach, the Acb must be capable of influencing several kinds of response. The Acb appears to mediate the motivational influence of noncontingent Pavlovian CSs on instrumental and locomotor behaviour — an influence that has been termed *incentive salience* (Robinson & Berridge, 1993; Berridge & Robinson, 1998), or Pavlovian incentive value (Dickinson *et al.*, 2000). One of the greatest remaining problems of Acb function is to understand the manner in which information passing through the Acb is encoded, and modified by this Pavlovian influence, and how the AcbC and AcbSh interact — apparently in different ways for different tasks — to provide this motivation.

Chapter 5.

Local analysis of behaviour in the adjusting-delay task for assessing choice of delayed reinforcement

Abstract. The adjusting-delay task introduced by Mazur (1987) has been widely used to study choice of delayed reinforcers. The adjusting-delay paradigm involves repeated choice between one reinforcer delivered after a fixed delay and another, typically larger, reinforcer delivered after a variable delay; the variable delay is adjusted depending on the subject's choice until an equilibrium point is reached at which the subject is indifferent between the two alternatives. Rats were trained on a version of this task and their behaviour was examined to determine the nature of their sensitivity to the adjusting delay; these analyses included the use of a cross-correlational technique. No clear evidence of sensitivity to the adjusting delay was found. A number of decision rules, some sensitive to the adjusting delay and some not, were simulated to examine which effects usually supposed to be a consequence of delay sensitivity could be explained by delay-independent processes, such as a consistent, unchanging preference for one of the alternatives.

INTRODUCTION

While delayed reinforcement can have profound effects on learning (e.g. Grice, 1948; Dickinson *et al.*, 1992), it can also affect choice behaviour in well-trained animals. The effects of delays to reinforcement on choice have been extensively investigated in the consideration of 'impulsive choice' (Ainslie, 1975), exemplified by the inability of an individual to choose a large delayed reward in preference to a small immediate reward. As discussed in Chapter 1 (p. 60), choice with delayed reinforcement may be assessed using free-operant tasks such as the concurrent-chains procedure (Davison, 1987) or in discrete trials. Discrete-trial tasks may be further subdivided into 'systematic' tasks (e.g. Evenden & Ryan, 1996), in which the experimenter varies the delay to one or more of several reinforcers and then measures choice, and 'adjusting' tasks, in which the subject's behaviour determines which delays are to be sampled.

The adjusting-delay task was introduced by Mazur (1984; 1987; 1988). Its principle is as follows. Subjects are given repeated choices of a small reinforcer A delivered after a small fixed delay (d_A , which may be zero to give immediate delivery) and a large reinforcer B delivered after a longer delay (d_B). The delay d_B may be altered; it is known as the *adjusting delay*. There is a rule for adjusting d_B depending on the subject's choices: if the subject consistently chooses the small ('fixed', 'unadjusting') reinforcer, the delay to the large reinforcer is reduced, while if the subject prefers the large reinforcer, the adjusting delay is increased. (It is assumed that subjects are sensitive to the changes in the adjusting delay.) The objective is that the adjusting delay tends to an equilibrium value d_B' , the 'indifference point' at which the effect of the delay of reward B cancels the effect of the larger magnitude of the reward and the two levers are chosen equally often. In practice, trials are usually grouped into blocks of four. The first two trials are forced presentations of each alternative separately, to ensure that the subject samples the currently pro-

grammed delays and reinforcers. The other two are free-choice trials. If the subject chooses the same alternative on both of these trials, the delay dB is altered according to the rules stated above. If the subject chooses each alternative once, dB is not altered. Subjects perform this task until dB has reached a stable value (various definitions of stability have been used) and the mean value of dB for stable trials is taken as dB'.

This task has provided strong support for the view that the effects of delayed reward are well described by a hyperbolic discount function (Mazur, 1987), and has been used with success in describing subjects' choice with delayed, probabilistic, and conditioned reinforcement (reviewed by Mazur, 1997). The effects of motivational and neurochemical manipulations have been clarified using this task (Wogar *et al.*, 1992; 1993b) and a version in which the magnitude of the reward is varied according to the same principles has also proved useful (1997a; Richards *et al.*, 1997b; 1999).

So far, this success has been on the 'molar' timescale; that is, based on values of dB' that are the mean of dB over a long series of choices on the part of the subject. The present study was designed to investigate choice behaviour in this task at a 'molecular' (trial-by-trial) level by examining the relationship between dB and choice, and to see if the task was suitable for neurotoxic lesion studies and acute pharmacological studies of impulsive choice. Rats were trained on the adjusting-delay task and their choices analysed to determine their sensitivity to dB. As a simple relationship between dB and choice was not found, computer simulations were conducted to investigate which observed features of performance can be explained by factors independent of dB, and evidence was sought of rats' sensitivity to their history of recent delays.

EXPERIMENT

Methods

Subjects

Eight experimentally naïve male Lister hooded rats were housed in pairs, provided with free access to water and were maintained throughout the experiment at 85% of their free-feeding mass. Housing conditions were described in detail in Chapter 2.

Adjusting-delay technique for assessing choice with delayed reinforcement

This behavioural task was based on those of Mazur (1987; 1988) and Wogar *et al.* (1992; 1993b). Four of the standard operant chambers were used (see Chapter 2), except that they were not fitted with dippers or traylights. The reinforcers used were 45-mg sucrose pellets (Rodent Diet Formula P, Noyes, Lancaster, NH). The apparatus was controlled by software written by R.D. Rogers, N. Daw and R.N. Cardinal.

Rats were first trained to press both levers (FR1 schedule, one-pellet reinforcer) in 30-min sessions daily, until a criterion of 50 presses per session was reached. The two levers were designated Levers A and B, counterbalanced left/right across subjects. Lever A produced immediate small rewards (1 pellet), while Lever B produced delayed larger rewards (2 pellets).

At the start of a session the houselight was switched on, and remained on for the duration of the session. Each session contained 10 trial blocks. Each block consisted of four lever presentations. The first two were forced-choice situations, with Levers A and B presented singly; the A/B order was randomized. Following these, there were two open-choice presentations of both levers simultaneously. Every presentation began with the illumination of the central magazine light, and the levers were extended 10 s later.

When the rat responded on a lever, the light above that lever was switched on, the magazine light was extinguished, and the levers were retracted. When the rat responded on Lever A, one pellet was delivered immediately. When it pressed Lever B, a delay ensued, after which two pellets were delivered. In both cases, the lever light was switched off as pellet delivery commenced. If the rat did not respond on a lever after a 'limited hold' period of 10 s, an omission was scored: the magazine light was switched off and the levers were retracted. No extra presentations were given to make up for omissions, but omissions were a very infrequent event (see *Results*).

If Lever A was chosen on both open-choice presentations, the delay associated with Reward B was decreased by 30% for the next trial block. If Lever B was chosen on both presentations, the delay was increased by 30%. If each lever was chosen once, the delay was not altered. The delay was initially 2 s and was kept within the range 2–20 s; this range was increased to 2–45 s from session 21 (trial block 201) as it became apparent that some rats had reached the maximum delay. From session 64, the delay was altered by 20% rather than 30%.

So that the choice of lever could not affect the frequency of reinforcer delivery, the time between lever presentations was kept constant at 45 s (or 70 s, after the maximum delay was increased). There were 10 trial blocks (of 4 lever presentations) per session, for a session length of 30 min (or 47 min after the increase). Adjusting delays for each subject were carried over from one session to the next as if there were no break.

Subjects were trained on this task for 80 sessions with one session per day.

Analysis of behavioural data

Choice-by-delay graph. To determine whether the current adjusting delay actually influenced the rats' choice, choice-by-delay plots were constructed. To create these plots, omissions were first excluded. Next, subjects' responses were assigned to a bin based on the adjusting delay that was operative at the time the response was made. (As the adjusting delay was altered on a logarithmic scale, bins of $0.1 \log_{10}$ units were used, though none of the analytical techniques used assumed that time was subjectively perceived by the subjects on a logarithmic scale.) For each rat, in each bin, a preference score was then calculated as the proportion of choices in which the delayed reward was selected.

To supplement this analysis, a simple measure of each rat's sensitivity to delay was derived by calculating the correlation between the rat's choice (Unadjusted lever, scored as 0, or Adjusted lever, scored as 1) and the logarithm of the adjusting delay was calculated. All data from every choice trial (except omission trials) were used, giving up to 1600 data for most rats. As the preference variable was dichotomous, the point-biserial correlation r_{pb} was calculated. This is numerically identical to the Pearson product-moment correlation coefficient r , and may be tested for significance in the same way (Howell, 1997, pp. 257/279–283). Once a correlation coefficient had been computed for each rat, the group's coefficients were compared to zero using a two-tailed t test to establish whether the group exhibited sensitivity to the adjusting delay.

Cross-correlations of preference and adjusting delay. In an attempt to elucidate the causal relationships between the adjusting delay and subjects' preference, cross-correlations were computed. Each rat's complete data set was examined using non-overlapping 'windows' of 10 choice trials (examining choice trials 1–10, then trials 11–20, and so on). Within each window, the preference for the adjusting alternative was calculated as the proportion of choice trials on which the adjusting alternative was chosen. For the same window, the mean \log_{10} (adjusting delay) was also computed. The calculated preferences and the mean adjusting delays were placed in temporal order to form two time series, and the cross-correlation function (CCF) of the two time series was computed. (Essentially, a cross-correlation computes the correlation between two functions at different lags and leads.)

This analysis attempts to separate out the influence of preference on delay from the influence of delay on preference, establishing the direction of causality. As preference was programmed to affect the adjusting delay in this task, it was expected that delays would be *positively* correlated with preference scores from the recent past (because preference was scored from 0, being exclusive preference for the unadjusting alternative, to 1, being exclusive preference for the adjusting alternative). Similarly, if long delays were aversive to the subjects, as might be anticipated, it

was expected that preference would be *negatively* correlated with delays from the recent past (equivalently, that delays would be negatively correlated with preference in the immediate future).

Mathematical background and pre-processing of data. Cross-correlation, a method within the discipline of time series analysis, depends upon a number of assumptions (see McCleary & Hay, 1980, pp. 229–273; Gottman, 1981, pp. 321–322). The topic is extremely complex and a thorough treatment will not be presented here. However, interpreting a CCF requires that both variables be ‘stationary’ — loosely, that there be no *autocorrelation* in either variable. (A variable exhibits autocorrelation, or is ‘non-stationary’, when its value at some time point can be predicted from the value of the same variable at a different time; a variable that does not exhibit this property is said to be stationary, or ‘white noise’.) Autocorrelation in either variable can introduce spurious correlation into the CCF; thus, a cross-correlation of autocorrelated variables is uninterpretable (McCleary & Hay, 1980, pp. 243–246). To correct for this, transformations are conducted before cross-correlating; this process is called ‘prewhitening’ and is performed on each variable separately, termed ‘double prewhitening’ (see also Hare, 1996, chapter 1). An example of this technique is given by Bautista *et al.* (1992).

To prewhiten a time series, an autoregressive integrated moving average (ARIMA) technique was used (Box & Jenkins, 1970; McCleary & Hay, 1980, p. 18; Gottman, 1981; see also StatSoft, 1999). Again, this will not be described thoroughly here, but the essence is to build a mathematical model of a time series that describes the autocorrelation in the time series, then to subtract the model’s predictions from the original data, removing the autocorrelation from the time series. Briefly, the notation ‘ARIMA(p,d,q)’ describes a mathematical model of a time series, specifying the degrees to which a time-lagged value of the variable is used as a predictor (autoregression; p), the number of passes on which the variable should be subtracted from a time-lagged version of itself before being used as a predictor (differencing; d), and the number of moving average parameters (q). As an example, an ARIMA(2,1,0) model contains two autoregressive parameters and no moving average parameters, calculated after the series has been differenced once. An autocorrelation function (ACF), which correlates a function with a time-shifted version of itself, may be used to identify the ARIMA model likely to provide the best fit to the data in question. The autocorrelation functions of ARIMA models are characterized by a discrete number of spikes corresponding to the moving average part of the model, and damped exponentials and/or damped sine waves corresponding to the autoregressive part of the model. Autocorrelation and partial autocorrelation functions (ACF, PACF) were computed for each variable being prewhitened (i.e. the preference score time series and the adjusting-delay time series) and used to identify autoregressive and/or moving average terms (that is, particular values of p , d , and q) as described by McCleary & Hay (1980), minimizing the number of terms included in the model (though model parsimony was considered secondary to obtaining a good fit). This model was then fitted to the variable, checking that it provided a significant fit, and the *residuals* were examined. If the residuals exhibited no autocorrelation (were of a white noise type), then the objective had been achieved: the autocorrelation had been removed from the original variable, and those residuals were used for cross-correlation. This technique is sometimes referred to as ‘filtering’ the original time series through an ARIMA model. It should be noted that the process of fitting an ARIMA model is empirical; a model is fitted to each time series separately (there are two time series from each subject), with the sole objective of removing autocorrelation from that time series. Importantly, the preference score and adjusting-delay time series were prewhitened *independently* before cross-correlation.

Finally, as the usual assumptions of correlation also apply to cross-correlation, the variables entered into the cross-correlation were checked for normality using the Kolmogorov–Smirnov test and by inspection of Q–Q plots (which plot the quantiles of the variable against the quantiles of a normal distribution).

To summarize, the following steps were conducted for each subject:

1. Calculate windowed choice ratios and log(adjusting delay), to give two time series.
2. Generate and fit an appropriate ARIMA model to each time series.
3. Cross-correlate the residuals from the two fitted ARIMA models.

Window size. The results of cross-correlational analysis depend in part upon the ‘window’ size used — for example, large windows permit more accurate calculation of preference, but they also obscure rapid, high-frequency changes in the cross-correlation coefficient. Pilot analyses were conducted with window sizes of 5, 10, 20, and 40. Smaller window sizes were not used, to avoid the preference score approaching a dichotomy, which would have violated the assumptions of the analysis. In all cases, the maximally *significant* cross-correlations were observed with the minimum window size used (5 choice trials); that is, the ‘optimal’ window size for detecting a correlation did not vary across subjects. The prewhitened data subjected to cross-correlation approximated a normal distribution even with this small window. Furthermore, the use of windows larger than 5 did not, in general, alter the lag at which the maximum cross-correlation was observed. As would be expected, larger windows yielded larger numerical correlation coefficients, but also increased the width of the confidence interval (as a larger window reduces the number of windows being analysed). Therefore, a window size of 5 was used for all subsequent analyses. Cross-correlations were computed out to lags and leads of 200 choice trials (40 decision windows).

Results

One rat (subject C7) fell ill and ceased responding from session 72; subsequent data from this rat were discarded. Other than this, responding was reliable, with rats failing to press a lever on only 1.53% of presentations. The obtained adjusting delays for 80 sessions (800 trial blocks) are shown in Figure 56 and typical individual records are shown in Figure 57. It is apparent that although the mean of the group of subjects appears relatively stable in the range 10–15 s, values of dB for individual subjects varied widely across the permissible range (which was 2–20 s for the first 20 sessions, and 2–45 s for the remainder).




Figure 56. Group mean adjusting delay for 8 rats, displayed by trial block. Thick and thin lines show mean ± 1 SEM. The boundaries between sessions are not shown. The maximum permissible value of the adjusting delay is shown as a stepped line at the top of the figure; this maximum was increased after session 20 (trial block 200).




Figure 57. Individual records, for all trial blocks. The thin grey line shows the maximum permissible value of the adjusting delay, as in Figure 56.

Choice-by-delay plots. Choice-by-delay plots are shown in Figure 58. As preference scores were arbitrarily calculated such that 0 represents exclusive preference for the unadjusted alternative (lever A) and 1 represents exclusive preference for the adjusted alternative (lever B), the theoretically predicted result would be a line of negative slope, indicating reduced preference for the large reinforcer at long delays. The obtained curve is relatively flat, indicating no effect of delay. If anything, Figure 58 suggests that a number of subjects had *high* preferences for the delayed reward when the delay was longest, and low preferences when the delay was low. A plausible interpretation is that the rats had a tendency to repeat their last response at extremes of delay — for example, a subject pressing the unadjusted lever many times in succession will drive dB down to its minimum permissible value, after which the subject can accumulate ‘unadjusted’ responses at the minimum delay. As a group, the point-biserial correlations did not differ from zero ($t_7 = 1.599$, NS; these data are also plotted later in Figure 61C, p. 161).

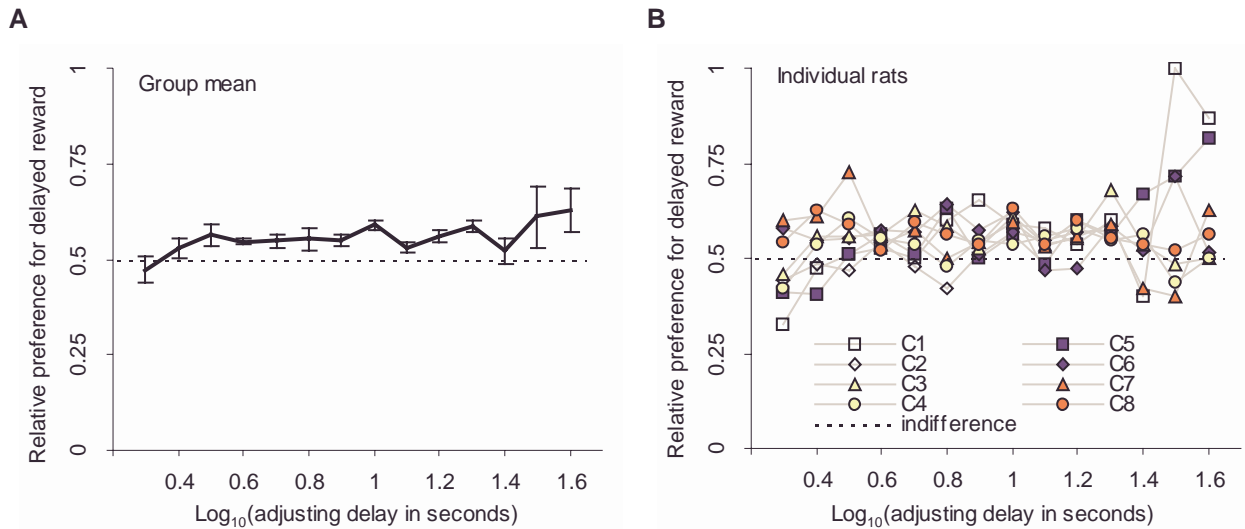



Figure 58. Choice-by-delay graphs for 8 rats. The ordinate (vertical axis) represents a preference score, from 0 (exclusive choice of the unadjusted, immediate lever) to 1 (exclusive choice of the adjusted, delayed lever), with omissions not analysed. The abscissa (horizontal axis) is $\log_{10}(\text{adjusting delay})$; preference was calculated in bins of 0.1 log units. **A:** Group mean \pm SEM. **B:** Individual subjects.

Slow changes in preference? It is probably unreasonable to expect rats to be perfectly sensitive to the adjusting delay currently in force. An obvious alternative is that the subjects are not immediately sensitive to changes in dB, despite the forced-choice trials, but rely on a slow cumulative learning process that gradually alters preference once the adjusting delay has been suboptimal for some time, leading to ‘overshooting’ and oscillation. (For example, a subject might prefer the large, delayed reinforcer when dB is low, leading to an increase in dB, yet fail to adjust its preference to reflect that increase for some time. By then, dB would have increased beyond the subject’s point of indifference, the small reinforcer would be preferred and the cycle would reverse. The value of dB would therefore oscillate around the indifference point rather than converging to it.) This is the view of several investigators (C.M. Bradshaw, personal communication, 7 October 1998; J.E. Mazur, personal communication, 16 November 1998). It may be termed a ‘running average’ hypothesis, since it suggests that the subjects are sensitive to some form of average of several recent values of dB.

If the ‘running average’ hypothesis is correct, then it is not so surprising that the choice-by-delay curve might be flat in its middle region. If dB oscillates around in the indifference point, there will be a range of values of dB for which the subject sometimes chooses the unadjusted alternative (at times when it is driving dB down), but sometimes chooses the adjusted alternative (when it is driving dB up). These tendencies might cancel out, leading to apparent indifference for this range of values of dB.

Cross-correlelograms. However, the ‘running average’ hypothesis predicts that choices should be correlated with adjusting delays from the recent past. Consequently, preference scores were cross-correlated with adjusting delays (as described in the *Methods*). For the prewhitening phase, it was found that an ARIMA(1,0,0) consistently described the vast majority of autocorrelation in the choice ratio time series, and an ARIMA(1,0,1) model was used for the delay time series. The final cross-correlations are plotted for each rat in Figure 59.

A



B




Figure 59. Cross-correlation of preference for the adjusting alternative with the adjusting delay. The analysis is organized so that positive lags indicate the effect of choice on delay, and negative lags indicate the effect of delay on choice. Thus, the ubiquitous positive correlation at small positive lags is the programmed rule for adjusting delay: when preference for the adjusting alternative is high, the delay is increased for later trials. A negative correlation at negative lags would indicate that high delays reduce subjects' subsequent preference for the delayed alternative. Confidence limits (horizontal dotted lines) are 2 SE. **A:** individual rats. **B:** group mean \pm SEM.

Figure 59 clearly reveals the contingencies programmed into the task: that choice affects delay, such that high preference for the adjusting alternative is strongly correlated with the adjusting delay in the near future. In contrast, it is not clear at all that the adjusting delay affected choice behaviour. (If the delay did

affect choice in the theoretically sensible direction, negative cross-correlations would be expected at negative lags in Figure 59.) Subjects' CCFs did exhibit occasional peaks in this region; to estimate the time course of the subjects' apparent sensitivity to the adjusting delay, the largest cross-correlation coefficients for each rat are printed in Table 15. It can be seen that there is considerable variation in the lag at which subjects' preferences were maximally correlated with the adjusting delay; some subjects' preferences appeared to be affected by the average adjusting delay from 15–40 choice trials previously (equivalent to 8–20 trial blocks, or 1–2 sessions), some subjects showed the maximal peak at 80–120 choice trials (or up to 6 sessions) previously. Additionally, none of these peaks is large. If the group is considered as a whole (Figure 59B), it is clear that no consistent sensitivity to past delays is seen.

Table 15. Maximum cross-correlations for each subject. The cross-correlations were computed using windows of 5 choice trials, and the correlation coefficient that was largest relative to its standard error at a negative lag is listed, as an index of the effect of delay upon the subject's choice. No attempt has been made to correct for multiple comparisons.

Rat	Lag (choice trials)	Correlation coefficient	Confidence limit (2 × standard error)
C1	–125	–0.141	–0.116
C2	–85	–0.107	–0.114
C3	–25	–0.123	–0.112
C4	–35	–0.125	–0.114
C5	–165	–0.146	–0.118
C6	–45	–0.127	–0.114
C7	–25	–0.15	–0.112
C8	–20	–0.173	–0.112

COMPUTER SIMULATIONS

It is intriguing that subjects performing a task that has produced highly consistent end-points in other studies (see Mazur, 1987; 1988) should show apparent insensitivity to dB. To establish what performance is possible using a decision rule that does not take account of the delay to reinforcement, a number of decision rules were simulated on a computer.

Methods

Computer simulations were written in the programming language C++ (Stroustrup, 1986; 2000); data from the simulations were fed into the same means of analysis as those from the real rats. Six decision rules were simulated, as follows:

Random. Decisions under the Random rule were independent of the adjusting delay. The adjusting alternative was chosen with probability 0.5 (and the unadjusting alternative also with probability 0.5).

Biased. The Biased rule was also delay-independent. The overall frequencies with which each rat chose the two alternatives were calculated (ignoring trials on which an omission occurred); each simulated subject was assigned the relative preference of one of the rats as its bias. On each choice trial, the adjusting alternative was selected with that probability, as shown in Table 16.

Biased-60. The Biased-60 rule implemented a fixed bias; the adjusting alternative was chosen with probability 0.6 (and the unadjusting alternative with probability 0.4).

Markov Chain. A Markov chain is an abstract entity that can be in one of several states at any given moment (phrased more obscurely, a 'finite state machine'). The chain is characterized by the set of probabilities of a transition occurring between each possible pair of states. In the present task, each choice alternative can be represented as a state (Adjusted and Unadjusted). A transition from the Adjusted state to the Unadjusted state would then represent a rat choosing the Adjusted lever on one choice trial, and the Unadjusted lever on the next trial.

Transition probabilities were calculated for each rat. After discarding omission trials, all choice trials were placed in order, and the relative frequency of the four possible transitions were computed. The transition probabilities are shown in Table 16. Eight Markov chain simulations were then performed, each simulation having the characteristic transition probabilities of one of the rats. The first choice made by each simulation was also the same as that of its corresponding rat.

Table 16. Overall proportion of choice responses on which the adjusting (Adj) alternative was chosen (used for the Biased rule), together with transition probabilities for each rat (used for the Markov chain decision rule). As omissions were ignored, pairs of transition probabilities sum to 1. The final column shows the first choice response ever made by each rat.

Rat	Overall proportion of Adj responses	$p(\text{Adj} \rightarrow \text{Adj})$	$p(\text{Adj} \rightarrow \text{Unadj})$	$p(\text{Unadj} \rightarrow \text{Adj})$	$p(\text{Unadj} \rightarrow \text{Unadj})$	First response
C1	0.470	0.544	0.456	0.406	0.594	Adj
C2	0.478	0.495	0.505	0.465	0.535	Adj
C3	0.570	0.630	0.370	0.496	0.504	Unadj
C4	0.530	0.579	0.421	0.476	0.524	Unadj
C5	0.489	0.573	0.427	0.409	0.591	Adj
C6	0.557	0.557	0.443	0.557	0.443	Adj
C7	0.551	0.554	0.446	0.546	0.454	Adj
C8	0.565	0.590	0.410	0.534	0.466	Unadj

Preference. The Preference rule was intended to mimic a ‘perfect’ subject — one whose choices immediately and accurately reflected the programmed adjusting delay. Each subject was assigned a preferred delay. On each choice trial, if the adjusting delay exceeded the preferred delay, the unadjusting alternative was chosen. If the adjusting delay was lower than the preferred delay, the adjusting alternative was chosen. If the adjusting delay exactly matched the preferred delay, the subject chose randomly ($p = 0.5$ for each alternative).

In order to match the Preference rule to the rats, each simulated subject was assigned a preferred delay derived from data from one rat; this preferred value was taken to be the mean adjusting delay over the last 200 trial blocks of testing (blocks 601–800). These values, in seconds, were 9.85 (subject C1), 3.55 (C2), 12.43 (C3), 9.20 (C4), 11.16 (C5), 10.83 (C6), 11.44 (C7: as this subject fell ill, its mean was calculated from trial blocks 601–710 only), and 11.50 (C8). The mean preferred delay for the simulations was thus 10.0 ± 0.99 s.

Running Average. The Running Average rule was also delay-dependent; it used the same basic decision rule as the Preference rule, and the preferred delays were calculated in the same manner. However, instead of comparing its preferred delay to the adjusting delay operative at that moment, the Running Average rule compared its preferred delay to the mean of dB over the last several choice trials. The actual decision window varied from subject to subject, and were chosen arbitrarily. The decision window sizes were drawn from a normally distributed random variable with a mean of 40 choice trials and an SD of 20 choice trials. The actual values used were 10, 60, 55, 52, 40, 22, 70, and 35 choice trials (mean 43, SEM 7).

It should be noted that the simulated Running Average rule has high mnemonic demands — the subjects remember every single delay within their decision window — and biologically more plausible algorithms exist (see Killeen, 1981, for a discussion), but it is a simple illustration of sensitivity to past delays that does not give heavy weighting to the most recent value. As more plausible algorithms often *do* give heavier weighting to more recent values (e.g. exponentially-weighted moving average; Killeen, 1981), the Running Average rule represents a stringent test of the analytical technique of cross-correlation as applied to this situation — if cross-correlation is observed with this rule, it would certainly be expected with more plausible algorithms.

Table 17 summarizes these decision rules.

Table 17. Summary of simulated decision rules. (Adj = selection of the adjusting alternative; Unadj = selection of the fixed alternative.)

Simulation name	Choice rules
<i>Delay-independent rules</i>	
Random	$p(\text{Adj}) = 0.5; p(\text{Unadj}) = 0.5$
Biased	The probability of selecting each alternative was fixed, and set to the overall probability with which one of the rats chose the alternatives (see Table 16).
Biased-60	$p(\text{Adj}) = 0.6; p(\text{Unadj}) = 0.4$
Markov chain	The probability of choosing each alternative was based solely on the previous choice, with the transition probabilities shown in Table 16.
<i>Delay-dependent rules</i>	
Preference	<ul style="list-style-type: none"> • if delay < preference, $p(\text{Adj}) = 1; p(\text{Unadj}) = 0$ • if delay > preference, $p(\text{Adj}) = 0; p(\text{Unadj}) = 1$ • if delay = preference, $p(\text{Adj}) = 0.5; p(\text{Unadj}) = 0.5$ Each subject had its own preferred delay, matched to one rat; these delays had a mean of 10.0 ± 0.99 s.
Running Average	Choice is determined as for the Preference rule, but the delay used to make the decision was the mean of dB over the 43 ± 7 most recent choice trials (see text).

For all decision rules, the starting conditions and the rules for updating the adjusting delay based on the subject's choice were identical to those in the real task (described earlier), including the change in the limits set on dB. Six decision rules were simulated, with 8 simulated subjects in each condition. Simulations were not repeated.

Application of a stability criterion to the Random decision rule. In a separate simulation, the Random rule was also used to establish the length of time needed for a randomly-deciding subject to meet the stability criteria previously used for pigeons by Mazur (1987; 1988; personal communication, 22 October 1998). The task simulated was changed so it matched exactly that used by Mazur (1988), as the starting value and stability criteria were unspecified in Mazur (1987). Thus, the starting adjusting delay was 8 s; trials were grouped into blocks of two single-lever trials and two choice trials, as before; the adjusting delay was altered arithmetically in steps of ± 1 s; the minimum value of dB was 0 s, and there was no maximum set on dB (as the time between reinforcement and the next trial was held constant in Mazur's study, rather than the time between the start of two consecutive trials). There were 64 trials per session (32 choice trials), the adjusting delay from one session was carried over directly into the next session, and subjects were tested for a minimum of 10 sessions. Data from the first two sessions were discarded and the rest of the data were tested for stability as follows. Each session was divided into two 32-trial blocks (i.e. 16 choice trials) and the mean adjusting delay for each block was calculated. The stability criteria were: (1) that neither the highest or the lowest single-block mean could occur in the last six blocks; (2) that the mean adjusting delay across the last six blocks was not the lowest or the highest such six-block mean; (3) that the mean of the last six blocks was within 10% of the mean of the preceding six (or within 1 s, whichever was greater). One hundred instances of the Random rule were simulated, and the time taken for each to meet these stability criteria was recorded.

Relationship between bias and mean adjusting delay. For a subject that takes no account of the adjusting delay, it is likely that the subject's bias has a systematic effect on the obtained mean adjusting delay, dB'. Firstly, to establish whether a manipulation that influenced a subject's bias could in principle affect dB', the mean value of dB was calculated over trial blocks 400–800 for each simulated subject using the Random or Biased-60 decision rules, yielding one value of dB' per subject (and $n = 8$ per group); these values were then subjected to a univariate ANOVA with the decision rule as a between-subjects factor.

Secondly, to establish the quantitative nature of the relationship between bias and dB', simulations were conducted using the conditions of Mazur (1988) described above. Each simulated subject was assigned a bias towards the adjusting lever; on every choice trial, it chose the adjusting lever with $p(\text{Adj}) = \text{bias}$, and $p(\text{Unadj}) = 1 - \text{bias}$. At every level of bias from 0.4 to 0.6 in steps of 0.01, one hundred subjects were simulated. The stability criteria described above were applied, and the mean value of dB over the last six (stable) half-session blocks was measured, just as in Mazur (1988, Table 1). This simulation was also repeated with dB limited within a range of 0–40 s.

Results

Local analysis of the simulated decision rules

The evolution of the adjusting delay is shown for the simulated decision rules in Figure 60 (compare the rat data in Figure 56). The Random rule simply generates a random walk between the limits set on dB. The Biased-60 rule chooses the adjusting alternative more frequently than the fixed alternative and thus drives the adjusting delay to high values, while wide excursions in dB are seen in the Random rule. The Preference rule generates tight oscillations around the preferred delay; even though the simulations' preferred delays were taken from the rat data, this simulation generates much less variability than the rats. The Running Average rule produces a sinusoidal oscillation around the preferred delay. As the group means shown in Figure 60 were derived from eight simulations, each with a different decision timebase, the group mean is not perfectly periodic (in fact, it represents a spectrum with eight frequency components). The delay-independent rules produced the pattern most like the rat data, with the Biased and Markov Chain rules generating values of dB in a similar range to the rats.

Figure 61 shows choice-by-delay plots for the simulated decision rules. This form of plot is clearly inadequate to demonstrate all but the simplest form of delay sensitivity: only the simple Preference rule demonstrates the theoretically predicted curve (high relative preference for the adjusting alternatives at low delays, and low preference at high delays). The other curves, including the delay-sensitive Running Average rule, are essentially flat. Indeed, the Running Average rule shows a *reduced* preference for the adjusting alternative at the minimum delay, probably due to repetition of responses — when the delay is high, for example, this rule begins to choose the Unadjusted lever and drives the delay to the minimum value; however, as its decisions are based upon several recent delays, it does not 'notice' that the delay has reduced for several trials, during which time it accumulates Unadjusted responses at the minimum delay.

Cross-correlational analyses of the decision rules were then conducted. The prewhitening process will be described first.

Prewhitening. For the Random, Biased, and Biased-60 simulations, as would be expected, there was no autocorrelation in the choice ratios but there was autocorrelation in the delays. As the adjusting-delay task only permits values of dB that are the same as, or a small way from, the value of dB in the preceding trial block, the delay at time t is correlated strongly with the delay at time $t - 1$; thus, this autocorrelation was modelled successfully by an ARIMA(0,1,0) model.

The same was true of the Markov Chain simulation. Even though each choice was programmed to depend (to a small extent) on the previous choice, the autocorrelation in choice ratios did not reach significance and no correction was made for it. The delay autocorrelation was again described by an ARIMA(0,1,0) model.

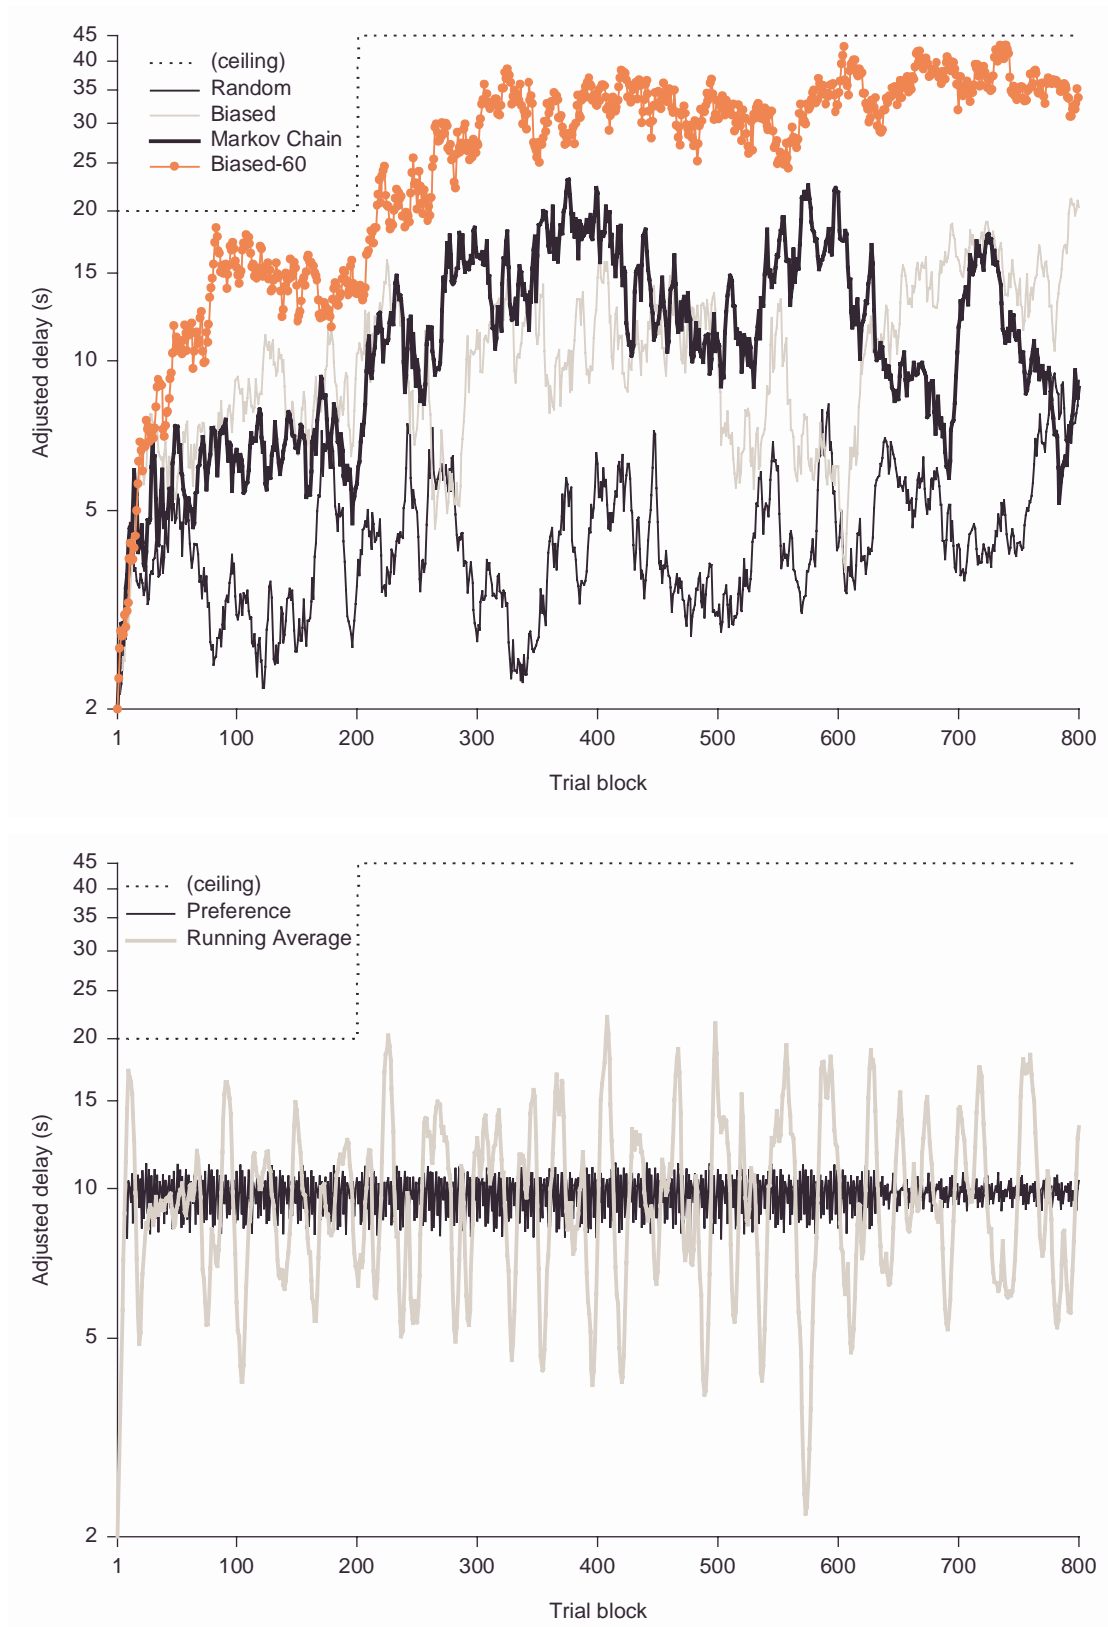


Figure 60. Group mean adjusting delay for the simulated decision rules, displayed by trial block. All simulations have $n = 8$. The top panel shows the delay-independent decision rules, and the bottom panel shows the delay-dependent rules.

The Preference simulation exhibited a significant autocorrelation of choice ratios (because its decisions at any one moment are closely related to decisions in the recent past, as the rule oscillates about its preferred value). This was removed by filtering the choice ratio time series through an ARIMA(2,0,1)

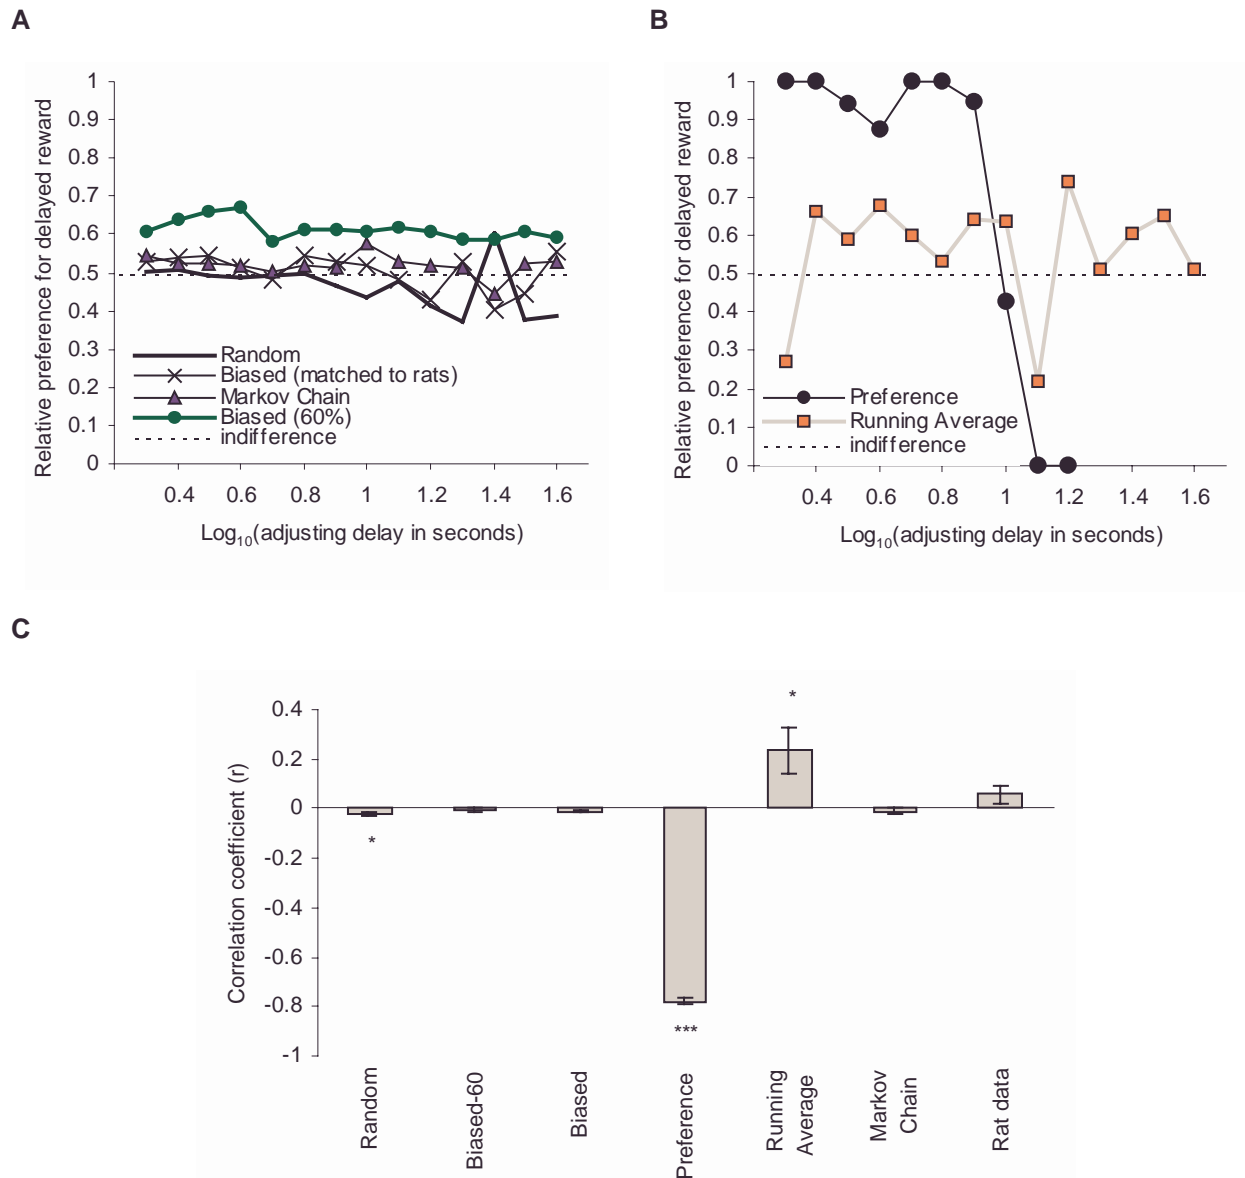


Figure 61. A & B: Choice-by-delay plots for the simulated decision rules. Only the simple Preference rule clearly demonstrates delay sensitivity in this plot, even though the Running Average rule is also delay-sensitive. **C:** Mean (\pm SEM) correlation coefficients for the correlation between preference for the adjusting lever and the adjusting delay. Correlation coefficients were calculated for each rat or simulated subject using data from all that individual's choice trials; the correlation coefficients were then compared to zero as a group using a two-tailed *t*-test (* $p < .05$; *** $p < .001$). The Preference rule exhibits consistent, negative correlation between preference and delay, indicating that it chooses the adjusting lever when the delay is low, and vice versa. The Running Average rule, which chooses on the basis of delays from the recent past, exhibits a small *positive* correlation between preference and the delay that is operating at the actual moment of choice. Ironically, the Random rule exhibits a significant (though very small) negative correlation! No other decision rule exhibited significant correlation; neither did the rats' choices.

model (determined following the method of Gottman, 1981, p. 262), and the delay time series through an ARIMA(4,0,3) model.

By its nature, the Running Average simulation makes decisions that are strongly correlated with decisions from the recent past. The 'recent' past in this case was quite long — the simulation with the 'longest memory' took account of delays from the last 70 choice trials (14 windows of 5 choice trials). Thus, a high-order ARIMA was necessary to capture the autocorrelation: an ARIMA(14,0,0) model was found to remove the vast majority of autocorrelation from both time series.




Figure 62. Cross-correlation functions for the simulated decision rules (group means \pm SEM), as in Figure 59. Confidence limits (horizontal dotted lines) are 2 SE. All decision rules exhibit a positive cross-correlation at positive lag, a result of the ubiquitous rule through which subjects' choices affect the adjusting delay. Only two rules (Preference and Running Average) exhibit a negative cross-correlation at negative lags, a phenomenon that suggests that the adjusting delay affects the subjects' choices, as indeed was the case for these and only these two rules.

CCFs. The cross-correlation functions for each decision rule are shown in Figure 62. Again, the CCF technique successfully detected the contingencies built into the task (the causal relationship: preference \rightarrow delay) in all cases. In addition, this technique successfully discriminated between rules that based their decisions upon the adjusting delay, and those that did not. Significant negative cross-correlations at nega-

tive lags (suggesting the causal chain: delay → preference) were detected for the Preference and Running Average rules, but for none of the delay-independent rules.

Inspection of individual records of the Random rule (Figure 63) revealed occasional ‘significant’ negative cross-correlations. As this decision rule was not influenced by the adjusting delay, there are two possible explanations. The first is failure of the prewhitening process to capture all of the autocorrelation in the delay time series; although prewhitening dramatically reduced the degree of autocorrelation, very small autocorrelations occasionally remained, having not been described by the ARIMA model. Autocorrelation can introduce spurious correlation into a CCF (McCleary & Hay, 1980). The second is simple statistical variation. The confidence intervals calculated by the statistical software used take into account the number of data points used to calculate the CCF, but not the number of leads and lags over which the CCF is computed and the number of comparisons this implies. The occasional isolated ‘significant’ correlation may therefore reflect Type I error (false rejection of the null hypothesis).

The relevance of this discussion is in the comparison with Figure 59A (p. 155), the cross-correlation data for the rats. It may be that occasional negative correlations observed in the rat data are due to the same processes that contributed to correlations in Figure 63, the Random simulations. While it remains possible that the rats exhibited genuine sensitivity to delay, though very slight and with a great deal of variation in its timescale across rats, the rats exhibited no evidence of systematic, consistent sensitivity to the adjusting delay, which is best assessed by consideration of the group mean (Figure 59B).




Figure 63. Cross-correlation functions for the simulated Random rule, plotted for each individual simulated subject (thin grey lines), together with the group mean (thick black line). Compare Figure 59.

Achievement of stability criteria by a delay-independent decision rule

Under the task conditions and stability criteria used by Mazur (1988), randomly-deciding simulated subjects reached stability after a mean of 15 sessions (range 10–43, SD 6, with 10 being the minimum number of sessions permitted by the criteria).

For comparison, Mazur (1988) found that pigeons reached stability in a mean of 14 sessions (range 10–29, SD 4; data taken Table 1 of Mazur, 1988, using all 61 conditions experienced by the four pigeons in which two choice trials were given per trial block). Figure 64 illustrates the point at which pigeons in Mazur’s (1987) experiment were considered stable.




Figure 64. Sample individual dB records of four pigeons, from Mazur (1987). Vertical bars mark the point beyond which performance was considered stable by Mazur's criteria.

Effect of bias on dB' using a delay-independent decision rule

When quasi-stable values of dB' from the Random and Biased-60 decision rules were compared (see *Methods*), it was found that the Biased-60 rule led to significantly higher values of dB' than the Random rule ($F_{1,14} = 649, p < 0.001$).

Results of the simulations designed to establish the relationship between bias and dB' are shown in Figure 65. When no limits were placed on dB, biasing the simulated subjects towards the adjusting lever increased the quasi-stable value of dB' (Figure 65A/C), the number of sessions to meet the stability criteria (Figure 65B), and the variance of these two measures. When the maximum value of dB was limited, manipulations of bias produced a sigmoid change in dB' (Figure 65D) with only minor effects on the number of sessions to criterion, which followed an inverted-U-shaped curve (Figure 65E).

There are regions of the curves in Figure 65A (magnified in Figure 65C) and Figure 65D that are approximately linear. Thus, if a subject chooses between the two levers in a way that is independent of dB, and a manipulation — such as a change in the reinforcement available on the unadjusting lever — were to affect its overall preference for the two levers, the obtained values of dB' might vary linearly with that preference (at least within the approximate range of preference of 0.45 to 0.55).

DISCUSSION

The present experiment failed to demonstrate that rats are sensitive to the rapidly-adjusting delay to reinforcement used in the task of Mazur (1987). The simulations suggested that even in the absence of such sensitivity, manipulations that affect subjects' overall preference for the two alternatives may have systematic effects on dB', the primary behavioural measure in this task. Furthermore, individual subjects did not, in general, exhibit stable patterns of choice (Figure 57, p. 153), a reason to question whether the task would be suitable for studying the effects of acutely-administered drugs on preference for delayed reinforcement. The simulations also indicated that the stability criteria previously applied to this task do not provide a guarantee that subjects are choosing other than at chance.

Interpretation of cross-correlational analysis

A number of analytical techniques were applied to this task for the first time. Analysis of the computer simulations demonstrated that correlating subjects' choices with the adjusting delay (dB) operative at the moment of choice successfully detects 'perfect' sensitivity to the adjusting delay (Figure 61), but fails to detect more complex forms of delay sensitivity such as sensitivity to a running average. The cross-correlational technique was more powerful, and successfully detected all the causal relationships embedded in the task itself (influences of choice on delay) and in the simulated decision rules (influences of delay upon choice, where applicable). When group data were considered, the cross-correlational analysis did not falsely detect causal relationships that were not present. This suggests that this technique, although complex, might be a useful way to understand the causal relationships operating in this schedule.

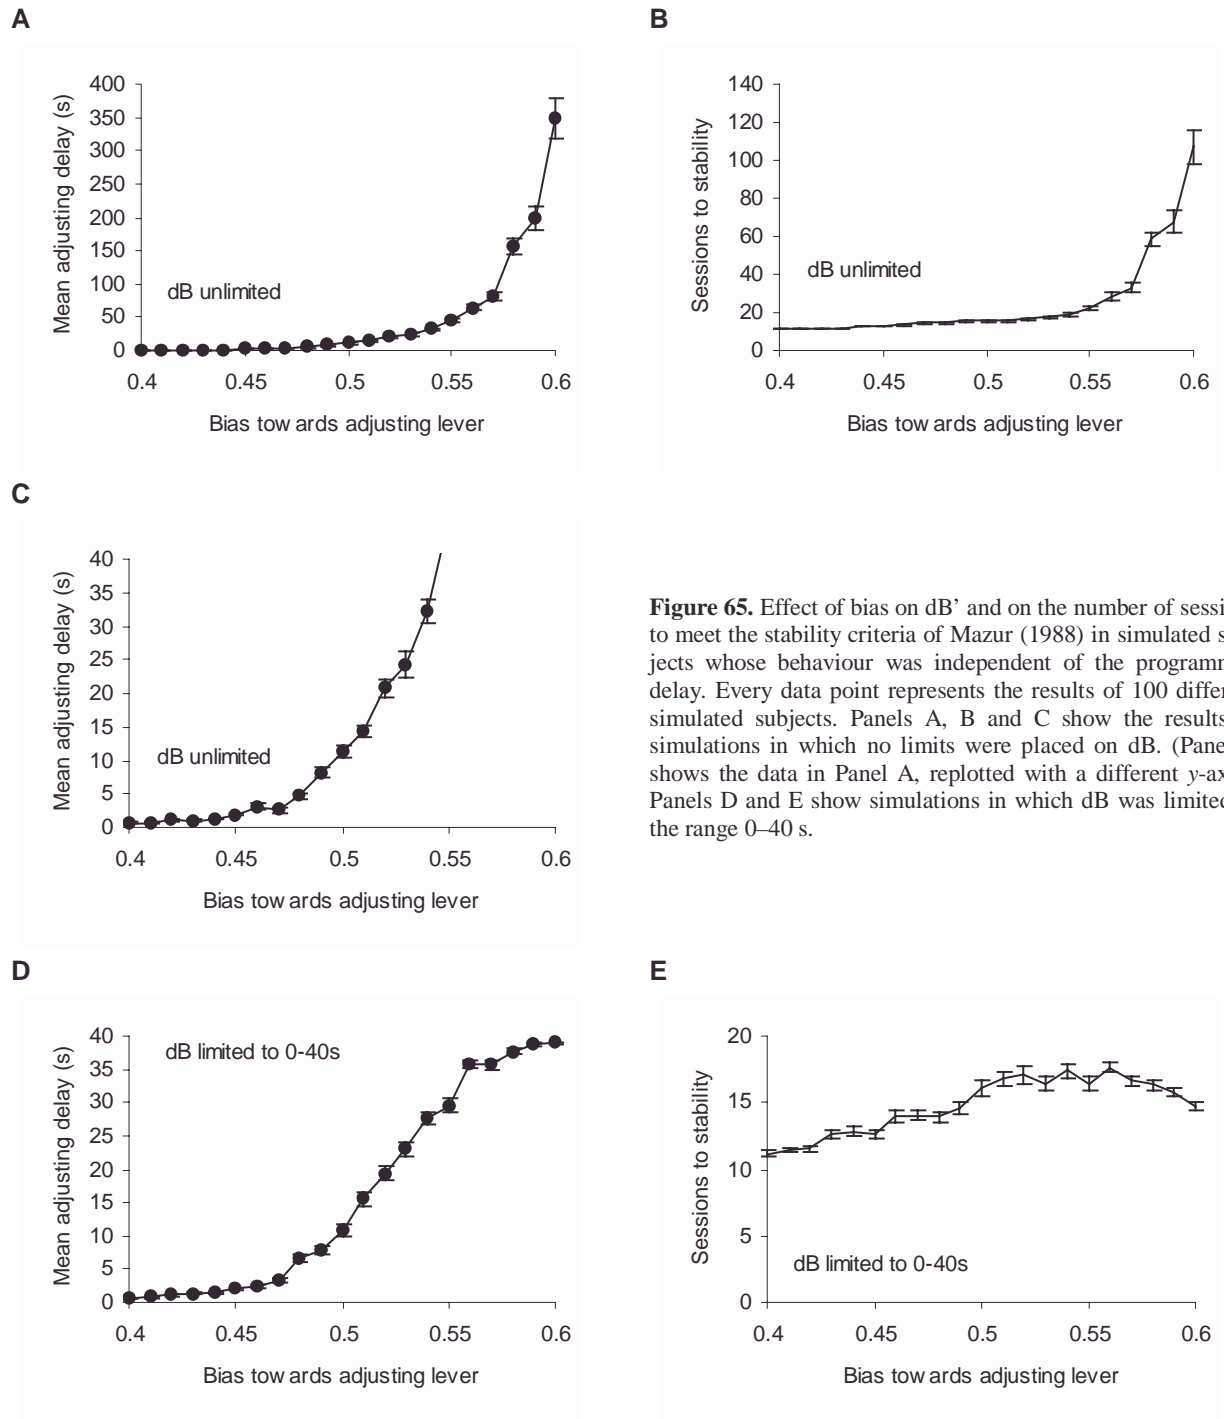


Figure 65. Effect of bias on dB' and on the number of sessions to meet the stability criteria of Mazur (1988) in simulated subjects whose behaviour was independent of the programmed delay. Every data point represents the results of 100 different simulated subjects. Panels A, B and C show the results of simulations in which no limits were placed on dB . (Panel C shows the data in Panel A, replotted with a different y-axis.) Panels D and E show simulations in which dB was limited to the range 0–40 s.

The technique failed to detect any consistent effect of the adjusting delay on the choices of the rat subjects.

It is not clear whether additional useful information can be gleaned from the prewhitening procedure that was applied to the data before it satisfied the assumptions of cross-correlation. For example, significant autocorrelation was only detected in the choice pattern of simulated decision rules that exhibited sensitivity to dB . (Autocorrelation was always observed in the sequence of values of dB , a consequence of the rules of the task.) Autocorrelation was, on the other hand, detected in the choice patterns of the rats, even though no influence of the dB upon choice was detectable by cross-correlation. This suggests that the rats exhibited a degree of cyclic behaviour that was unrelated to the adjusting delay. The suggestion of

cyclicity is borne out by inspection of individual records in Figure 57 (p. 152); the surprising finding is that this cyclicity is apparently not a direct consequence of the adjusting delay.

Stability does not imply sensitivity to the adjusting delay

An important point that emerges from these simulations is that apparent stability cannot be taken as evidence of subjects' titrating their preference between the two alternatives. For example, some of the data series shown in Figure 60 could be taken as stable by visual inspection. Even when investigators use formalized stability criteria, delay sensitivity is not implied. As discussed on p. 158, Mazur (e.g. 1987; 1988) has used quite strict criteria to determine when a subject has reached stable performance. In addition, Mazur reduced the likelihood of finding spurious stability by taking the final value of dB from one session as the starting value for the next, rather than applying the smoothing technique used by Wogar *et al.* (1992; 1993b), in which the mean value of dB for the last half of one session is taken as the new starting point. Nevertheless, randomly-deciding simulated subjects achieved Mazur's criteria within times comparable to real pigeons (p. 163).

Possible reasons for the present failure to observe sensitivity to dB

The first explanation that must be considered is that the rats *were* sensitive to dB, but in a way that was not detected by the present analyses; perhaps the sensitivity was fleeting, or its nature changed across the course of the experiment and was masked by analysing the entire sequence of choices made by each rat. Other than the occasional cross-correlational peaks that reached significance (Figure 59A, p. 155), which were also apparent in a delay-independent simulation (Figure 63, p. 163), no evidence was found for delay sensitivity in the rats. Still another possibility is that the rats did not generalize from the forced exemplar presentations to the choice trials, and thus their preference for the adjusting alternative depended upon how often they had sampled it recently, as well as upon dB, in a highly complex feedback manner.

A more obvious explanation is that the rats were not sensitive to dB at all. Of course, the present results may not be representative of performance on this schedule generally; the lack of sensitivity may have been a consequence of procedural differences between the present experiment and previous studies. These differences may be enumerated as follows:

1. The reinforcers used were one and two 45-mg sucrose pellets. As Mazur (1987) used 2 s or 6 s of access to grain as the reinforcer for pigeons, a larger relative magnitude, it may be argued that the rats in the present study failed to discriminate between the large and small reinforcers; different results might have been obtained if the delayed reinforcer had been larger. However, at least two rat-based studies of the adjusting-delay schedule have used one and two 45-mg food pellets as the reinforcers, with 'molar' behavioural results that indicated that the subjects discriminated between them (Wogar *et al.*, 1992; 1993b).
2. In the present experiment, the adjusting delay dB was varied by 30% at a time (or 20% for the last part of the experiment). In the original studies of pigeons (Mazur, 1987; 1988), dB was altered arithmetically, typically by ± 1 s; changing dB by 30% may have resulted in large swings in preference. However, proportional alterations of 30% have previously been used successfully (Wogar *et al.*, 1992; 1993b). Furthermore, Mazur (1988) has shown that increasing the step size has relatively little effect on the stable value of dB', though larger steps produce greater variability (as might be expected) and dB' sometimes increases with large step sizes.
3. The adjusting delay was not allowed to go below 2 s. In Mazur's early experiments, the floor on dB was zero; obviously, a zero floor is not possible with a proportional alteration, but it is true that

studies using proportional alterations (Wogar *et al.*, 1992; 1993b) have not placed a floor value on dB. It is possible, therefore, that the titration procedure failed because of this. If temporal discounting were steep enough that the fixed alternative (one pellet delivered after 0 s) was preferred to two pellets after 2 s, the indifference point would not be achievable, and subjects would simply keep dB at its floor value. It is possible that subjects C1, C2, and C5 attempted to do so (Figure 57, p. 153), though the CCF analysis did not demonstrate that they did, and the other five subjects certainly did not. In general, rats appear to be better able to wait for delayed reward than pigeons (see Mazur, 2000), making this interpretation less likely.

4. Similarly, a ceiling was placed on dB; initially, this was 20 s (following Wogar *et al.*, 1993b), though it was found necessary to increase this in the course of the experiment. The pigeon studies mentioned above did not place a ceiling on dB; however, as Figure 57 (p. 153) shows, once the ceiling was raised to 45 s, no rat preferred the adjusting alternative exclusively.
5. Trials were presented at constant intervals, in order to ensure that subjects could not do better by choosing 'small and often' instead of 'large but infrequent' rewards. This inevitably enforces a ceiling on dB. In several studies using this schedule (1987; Mazur, 1988), the time *between* trials was fixed; thus, the possibility existed for subjects to do well by choosing the smaller reinforcer and so being able to gain reward more often. Despite this procedural difference, systematic variations in the ITI do not appear to affect dB' in pigeons (Mazur, 1988).

Manipulations of dA and other parameters were not conducted in the present experiment, as a primary purpose of the experiment was to establish rapidly whether the task would be suitable for pharmacological and lesion studies. The present results are therefore not conclusive, as it has not been shown that the molar results of the present experiment are comparable to previous work (e.g. 1987; Mazur, 1988). In particular, it has not been demonstrated that dB' responds to long-term changes in dA in the same subjects that are insensitive to dB. However, this possibility will be explored briefly.

Effects of manipulations that alter subjects' preferences in a delay-independent manner

Effects of extrinsic manipulations

The present simulations show that evidence of an alteration in dB' as a result of a behavioural or neural manipulation is not proof of delay-dependent decision-making. Comparison of the group means from the Random and Biased-60 simulations (Figure 60, p. 160; analysis, p. 164) demonstrated that differences in relative preference for the two alternatives can lead to differences in dB', even though the decision rules generating these data took no account of dB. (This analysis illustrated a between-group difference, but the principle applies equally to a within-subjects manipulation.) Therefore, caution should be exercised when interpreting individual or group differences in dB' as an effect of a manipulation on delay sensitivity.

Effects of dA on dB'

It is clear that pigeons performing on an adjusting-delay schedule are sensitive to variations in the delay to reinforcement of the unadjusting alternative (dA) (e.g. Mazur, 1988; 1997); typical results are reproduced in Figure 66. On the basis of the present data, it is tentatively suggested that subjects performing this task are unable to track changes in dB. According to this hypothesis, they are unable to choose on the basis of the rapidly changing delay dB, and so come to assign a certain 'overall value' to the adjusting alternative. The perceived value of the unadjusting alternative, however, is constant over long periods of time, and when it changes suddenly, the 'value' assigned to the unadjusting lever changes accordingly.




Figure 66. Alterations in dA ('delay for 2-sec reinforcer') affect dB' ('delay for 6-sec reinforcer') in an orderly fashion in four pigeons. From Mazur (1987).

On any given choice trial, subjects ignore the current value of dB but instead compare the value of the unadjusting alternative with the 'overall' value of the adjusting alternative, giving rise to a dB-independent preference. The results of the simulations depicted in Figure 65 (p. 165) show that this relative preference may be translated into a quasi-stable value of dB', and that preferences within a certain range (approximately 45–55% preference for either alternative) are related near-linearly to the value of dB'. The rats in the present study made 52.6% of free-choice responses on the adjusted lever on average (range 47.0–56.5%, SD 4.1%), clearly in the range in which a manipulation affecting relative preference could alter dB'. In summary, this hypothesis states that subjects are sensitive to dA but not directly to dB. While this may not be an appealing idea, it seems possible.

Indeed, it has been observed that bias for the adjusting alternative, measured as the ratio of dB' to dA when the two reinforcers are equal ('bias for or against the adjusting procedure itself'; Mazur, 1984, p. 429), increases as a function of dA (Mazur, 1984, p. 431; though see Mazur, 1987, p. 63; Mazur, 1988, p. 46).

In principle, similar arguments regarding manipulation effects and bias apply to adjusting-magnitude tasks. The adjusting-magnitude task (Richards *et al.*, 1997b) is learned faster than the adjusting-delay task (see Ho *et al.*, 1999, p. 369), suggesting that rats may learn the contingencies more readily with varying reinforcer magnitudes than with varying delays to reinforcement (an interpretation compatible with studies of instrumental learning with delayed reinforcement, e.g. Lattal & Gleeson, 1990; Dickinson *et al.*, 1992). However, the adjusting-magnitude task also involves a titration method in which subjects' preference affects a variable that is assumed to affect subjects' preference in turn.

To emphasize a point, the present simulations do not prove that subjects in previous studies were insensitive to dB, but demonstrate that many of the observed molar features of performance can be obtained in the absence of such sensitivity.

Comparison to free-operant schedules of reinforcement

In contrast to the present experiment, rats are known to be able to track reinforcement rate extremely rapidly in some circumstances, and probably do so by timing interreinforcement intervals (IRIs). This was demonstrated by Mark & Gallistel (1994), who used two concurrent variable interval (VI) schedules of

lateral hypothalamic stimulation, ranging from VI 4 s to VI 256 s. They showed that the rats' response allocation tracked not only changes in the programmed ratio of reward between the two levers, but also the unprogrammed random fluctuations in the VI schedule, to an extent that their behaviour was governed by a very few of the most recent IRIs. This result implies that rats do *not* maintain and use a decaying 'running average' of the reward history, at least in that task (see Mark & Gallistel, 1994, pp. 90–91); Mark & Gallistel argue persuasively that their rats tracked the relative ratio of reward rate on the two levers by timing the interval to detect a fixed number of rewards (this number being from one to three).

It is an interesting question as to why rats are apparently capable of timing intervals on a seconds-to-minutes timescale and updating choice behaviour based on these intervals in concurrent VI schedules, but are apparently incapable of this in the discrete trials adjusting-delay procedure. It must be acknowledged that the two procedures are very different. Discrimination of changes in relative reinforcement rate may be easier than discrimination of changes in reinforcement delay in a discrete-trial procedure. One possibility, discussed by Mark & Gallistel (1994, p. 94) is that regular, dramatic changes in reward encourage extreme sensitivity to these changes, while relative stability with slow changes in reinforcement parameters (as in the present task) discourages local sensitivity to the reinforcement contingencies. Whether this reflects the operation of two psychological processes is unclear, but relative invariance of response–reinforcement contingencies has been suggested to be the key factor engendering habitual responding (Dickinson, 1985) (as discussed in Chapter 1, p. 25); discrete-trial schedules constrain behavioural variability much more than free-operant schedules. One highly speculative interpretation is that the task of Mark & Gallistel (1994) tests goal-directed action while choice in Mazur's (1987) procedure is more heavily influenced by the relative strength of two differentially reinforced stimulus–response habits.

SUMMARY

The adjusting-delay task has produced consistent results on the molar scale and lends itself well to using dB' values as a measure of relative preference of different 'fixed alternative' conditions (work reviewed by Mazur, 1997). However, caution must be exercised when interpreting effects on dB' as changes in sensitivity to dB. In the present study, rats did not update their behaviour rapidly to reflect changes in dB, and no clear evidence for any form of sensitivity to dB was found. These results suggest that rats' behaviour on this task would not be characterized well as 'informed choice'; the psychological mechanisms underlying choice in this task are not clear at present. Artificial decision rules that take no account of dB were found to be able to replicate a number of observed features of performance on the task, including the satisfaction of stability criteria and the generation of within- or between-subject differences in dB'. Finally, rats' preference did not exhibit clear stability or consistency even after prolonged training. The task therefore appears unsuitable for acute pharmacological studies, for which it would be preferable to be able to perturb and re-stabilize performance within one or a few sessions. In Chapter 6, a task of the 'systematic' kind will therefore be turned to, in which the subject has no influence on the delay to reinforcement.

Chapter 6.

The effects of *d*-amphetamine, chlordiazepoxide, alpha-flupenthixol and behavioural manipulations on choice of signalled and unsignalled delayed reinforcement

Abstract. Inability to tolerate delays to reward is an important component of impulsive behaviour, and has been suggested to reflect dysfunction of dopamine systems. The present experiments examined the effects of signalling a delayed, large reward on rats' ability to choose it over a small, immediate reward, and on the response to amphetamine, a dopamine receptor antagonist, and a benzodiazepine. Three groups of Lister hooded rats were tested on a two-lever discrete-trial delayed reinforcement task in which they chose one pellet delivered immediately or four pellets delivered after a delay. This delay increased from 0 to 60 s during each session. Trials began with illumination of a houselight: in the Houselight group, this remained on during the delay and feeding period. In the No Cue group, the houselight was extinguished at the moment of choice. In the Cue group, a stimulus light was illuminated during the delay. Once trained, the rats were challenged with *d*-amphetamine (0.3, 1.0, 1.6 mg/kg), chlordiazepoxide (1.0, 3.2, 5.6, 10 mg/kg), α -flupenthixol (0.125, 0.25, 0.5 mg/kg), and various behavioural manipulations. Subjects' choice became and remained sensitive to the delay; the cue speeded learning. Amphetamine decreased choice of the large reinforcer in the No Cue group and increased it in the Cue group. α -Flupenthixol and chlordiazepoxide generally decreased preference for the delayed reinforcer; flupenthixol reduced the cue's effects, but chlordiazepoxide did not interact with the cue condition. It is concluded that signals present during a delay can enhance the ability of amphetamine to promote choice of delayed rewards.

INTRODUCTION

Among the many features of impulsivity, one is 'impulsive choice', exemplified by the inability of an individual to choose a large delayed reward in preference to a small immediate reward (Ainslie, 1975). Impulsive choice has been suggested to reflect an alteration in reinforcement processes, namely that delayed reinforcers have lost their effectiveness, and has been suggested to underlie attention-deficit/hyperactivity disorder (ADHD; Sagvolden *et al.*, 1998; Sagvolden & Sergeant, 1998). ADHD is amenable to treatment with psychomotor stimulant drugs (Bradley, 1937; see Solanto, 1998 for a recent review), suggesting that they might promote the choice of delayed rewards. However, in laboratory models of impulsive choice, the effects of acute administration of psychostimulants have varied: some studies have found that they promote choice of delayed reinforcers (Sagvolden *et al.*, 1992; Richards *et al.*, 1999; Wade *et al.*, 2000), while others have found the opposite effect (Charrier & Thiébot, 1996; Evenden &

Ryan, 1996), and it has been shown that the same psychostimulant can have opposite effects in different tasks designed to measure impulsivity (Richards *et al.*, 1997a).

In studies of delayed reinforcement, it has been demonstrated that signalled delays generally maintain higher rates of free-operant responding than unsignalled delays (see Lattal, 1987 for a review), and signals present during the delay can have an important role in discrete-trials choice (Mazur, 1997). A signal or cue that is associated selectively with a reinforcing outcome may become a conditioned reinforcer (Figure 67). Conditioned reinforcement can affect choice behaviour, perhaps the best demonstration being that of Williams and Dunn (1991), in which pigeons preferred a key associated with a conditioned reinforcer despite this leading to fewer presentations of food. Since amphetamine-like drugs potentiate the effects of conditioned reinforcers (Hill, 1970; Robbins, 1976; Robbins, 1978; Robbins *et al.*, 1983), amphetamine may promote choice of signalled delayed reinforcement.

Evenden and Ryan (1996) developed a model of impulsive choice in which food-restricted rats chose between a small, immediate reward and a large, delayed reward in discrete trials, the delay to the large reinforcer being increased in steps as the session progressed. The present study examined the effects of the psychostimulant *d*-amphetamine, the benzodiazepine chlordiazepoxide, and the mixed dopamine D₁/D₂ receptor antagonist α -flupenthixol on performance of a modified version of this task, with particular emphasis on the effects of a signal present during the delay to reinforcement. Subsequently, to characterize the basis of performance on the task, the effects of this signal itself, of removing the delays, reversing the order of the delays, of satiation, and of extinction were examined.

Three groups of animals were trained on variations of the task, differing only in the signalling conditions. In the Cue condition, illumination of a stimulus light during the delay provided a signal that was unambiguously associated with the large reinforcer only. This design is commonly used to establish stimuli as conditioned reinforcers in delay-of-reinforcement experiments (for reviews, see Williams, 1994a; Mazur, 1997). In the No Cue condition, the rats awaited and collected the reinforcers in darkness, with no signal present during the delay. This closely resembles the situation in Evenden and Ryan's (1996) study. The Houselight condition was intermediate between these: in this condition, the houselight was illuminated at the start of the trial and remained on until 6 s after the subject had collected the reward. The houselight therefore preceded and accompanied delivery of the large and small reinforcers.

Given that the effect of amphetamine on performance of this task in the absence of differential cues was to increase preference for the small immediate reward (reduced tolerance of delay, Evenden & Ryan, 1996), the addition of a conditioned reinforcer would be expected to reduce or reverse this effect. (The Houselight group were predicted to be intermediate or equivalent to the Cue group, in that the houselight is a weak predictor of food.) Chlordiazepoxide was used as a positive control; its effects were not expected to differ in the presence of a cue because benzodiazepines do not affect the action of appetitive conditioned reinforcers (Killcross *et al.*, 1997a), while the dopamine receptor antagonist α -flupenthixol was predicted to have opposite effects to amphetamine in the cue condition as it attenuates the effects of conditioned reinforcers (Robbins *et al.*, 1983; Killcross *et al.*, 1997a).

METHODS

Subjects, apparatus, and behavioural task

Subjects were 24 experimentally naïve male Lister hooded rats maintained at 90% of their free-feeding mass and housed in pairs (for details of housing conditions, see Chapter 2).




Figure 67. Choice of signalled and unsignalled delayed reinforcement. Subjects may choose between a small, immediate reward and a large, delayed reward. In the 'Cue' condition, a stimulus light is illuminated during the delay to reinforcement; this stimulus is therefore paired with the large reinforcer and may become a conditioned reinforcer.

Systematic technique for assessment of preference for delayed reinforcement

The standard operant chambers described in Chapter 2 were used, with 45-mg sucrose pellets (Rodent Diet Formula P, Noyes, Lancaster, NH) as the reinforcer.

Training. Subjects were first trained under an FR1 schedule to a criterion of 50 presses in 30 min, first for the left lever and then for the right. They were then trained on a simplified version of the full task. The session began with the levers retracted and the operant chamber in darkness. Every 40 s, a trial began with illumination of the houselight and the traylight. The subject was required to make a nosepoke response within 10 s, or the current trial was aborted and the chamber returned to darkness. If the subject nosepoked within this time limit, the traylight was extinguished and a single lever presented. If the rat failed to respond on the lever within 10 s, the lever was retracted and the chamber darkened, but if it responded, a single pellet was delivered immediately and the traylight was illuminated until the rat collected the pellet (or a 10-s collection time limit elapsed, whereupon the chamber was darkened). In the Houselight condition, the houselight was left on until 6 s after the food had been collected; in the Cue and No Cue conditions it was switched off at the moment the lever was pressed.

In every pair of trials, the left lever was presented once and the right lever once, though the order within the pair of trials was random. Rats were trained to a criterion of 60 successful trials in one hour (the maximum possible with a 40-s period being 90).

Behavioural procedure. The task was based on Evenden and Ryan's (1996) procedure and is illustrated in Figure 68. Aside from the use of an extra signal during the delay, the present task differs from that of Evenden and Ryan in a number of ways; in particular, the subjects were required to initiate the trials and choose a lever within a limited time, and a forced-choice trial on each lever was given at the start of each block of choice trials at a given delay. Additionally, in their procedure the houselight was always on, whereas in the present studies the houselight was extinguished during the intertrial interval (ITI), making it an informative stimulus (in that food was delivered




Figure 68. Schematic of the task. On the *right-hand side*, the format of a single trial is shown. This diagram shows in detail the Houselight condition, in which the houselight remains on from the start of the trial until 6 s after the subject has collected the reward. On the *left-hand side*, the differences between the three lighting conditions are illustrated. In the No Cue condition, the houselight is switched off at the moment of choice. In the Cue condition, the houselight is similarly switched off when the subject responds on a lever, but a stimulus light is illuminated during the delay that precedes delivery of the large reinforcer.

when the houselight was on, but never when it was off). Finally, subjects were not given exposure to the large reinforcer before delays were introduced into the task.

The session began in darkness with the levers retracted; this was designated the intertrial state. Trials began at 100-s intervals. Each trial began with the illumination of the houselight and the traylight. The rat was required to make a nosepoke response, ensuring that it was centrally located at the start of the trial (latency to poke was designated the initiation latency). If the rat did not respond within 10 s of the start of the trial, the operant chamber was reset to the intertrial state until the next trial began and the trial was scored as an omission. If the rat was already nosepoking when the trial began, the next stage followed immediately.

Upon a successful nosepoke, the traylight was extinguished and one or both levers were extended. One lever was designated the Delayed lever, the other the Immediate lever (counterbalanced left/right). The latency to choose a lever was recorded. (If the rat did not respond within 10 s of lever presentation, the chamber was reset to the intertrial state until the next trial and the trial was scored as an omission.) When a lever was chosen, both levers were

retracted. Choice of the Immediate lever caused the immediate delivery of one pellet; choice of the Delayed lever caused the delivery of 4 pellets following a delay. In the Cue condition, the houselight was switched off at the moment of choice and a stimulus light above the chosen lever switched on for the duration of the delay. In the No Cue condition, the stimulus light was not switched on. In the Houselight condition, the houselight remained on instead. These three conditions are illustrated in Figure 68.

Following any delay, the stimulus light was switched off, the traylight was switched on and the reinforcer for that lever was delivered. Multiple pellets were delivered 0.5 s apart. If the rat collected the pellets before the next trial began, then the time from delivery of the first pellet until a nosepoke occurred was recorded as the collection latency. The traylight was switched off, and in the Houselight condition the houselight remained on for another 6 s (eating time). In other conditions there was no houselight illumination during this time. If the rat did not collect the food within 10 s of its delivery, the operant chamber entered the intertrial state, though collection latencies were still recorded up to the start of the next trial. The chamber was then in the intertrial state and remained so until the next trial. There was no mechanism to remove uneaten pellets, but failure to collect the reward was an extremely rare event (see *Results*).

The delay was varied systematically across the session. A session consisted of 5 blocks, each comprising two trials on which only one lever was presented (one trial for each lever, in randomized order) followed by ten free-choice trials. Delays for each block were 0, 10, 20, 40 and 60 s respectively. As trials began every 100 s, the total session length was 100 minutes; subjects received one session per day.

Pharmacological and behavioural manipulations

A stability criterion was defined as follows: after excluding single-lever trials, choice ratios (delayed lever responses ÷ total responses) were calculated for each rat using the summed responses for three consecutive sessions, and subjected to analysis of variance (ANOVA) with delay as a within-subjects factor. When the effect of delay was significant at the $\alpha = .01$ level, the rats were considered to have criterion performance from the first session of the three. The degree of sensitivity to the effects of the delay within each session was also assessed by calculating the slope of the linear regression of %choice of the large reinforcer against $\log(\text{delay} + 1 \text{ s})$ for each subject, though this measure did not form part of the criterion. Following attainment of the criterion, baseline assessments were performed on seven sessions immediately prior to the start of pharmacological and behavioural manipulations, which were conducted as listed in Table 18.

Table 18: Experiments performed.

Group	Houselight during delay and feeding period	Stimulus light during delay	Manipulations, in order
Houselight ($n = 8$)	On	Off	amphetamine 1.0, 0.3 and 1.6 mg/kg omission of delays addition of cues hungry versus sated (rapid) hungry versus sated (longer term) descending delays
No Cue ($n = 8$)	Off	Off	amphetamine 1.0, 0.3 and 1.6 mg/kg chlordiazepoxide 10.0, 1, 3.2 and 5.6 mg/kg α -flupenthixol 0.25, 0.125 and 0.5 mg/kg extinction
Cue ($n = 8$)	Off	On	amphetamine 1.0, 0.3 and 1.6 mg/kg chlordiazepoxide 10.0, 1, 3.2 and 5.6 mg/kg α -flupenthixol 0.25, 0.125 and 0.5 mg/kg omission of cues

Drugs. *d*-Amphetamine sulphate, α -flupenthixol dihydrochloride and chlordiazepoxide hydrochloride (Sigma, UK) were all dissolved in sterile 0.9% saline to give a final volume of 1 ml/kg and injected intraperitoneally 10 min before the start of the session (60 min for flupenthixol). Doses were calculated as the salt and are listed in Table 18.

Drug studies. Each dose was tested over six sessions, with each rat experiencing either DVDVDV or VDVDVD (D drug session, V vehicle), counterbalanced across rats. Responding under each dose was compared with responding during the vehicle sessions that alternated with that dose. This approach has the advantage of being able to compare each drug dose with vehicle data collected across the same time period, increasing the power to detect drug effects if the baseline shifts gradually; it also implies that any drug carry-over effects would reduce the power to detect effects. Collecting data for three drug and three vehicle sessions enabled accurate determination of choice by giving 30 choice trials at each delay/dose combination. Between each six-session dose study, at least two days elapsed on which no injections were given.

Omission of delays. Following testing with amphetamine, the Houselight group were not included in further pharmacological studies but were tested under a range of behavioural manipulations. To establish whether they were still sensitive to the delays, they were first tested on six sessions alternating between the normal task and a version in which all delays were zero. Half of the rats began with the Delay and half with the No Delay condition.

Introduction of a cue. The Houselight group were next tested with successive sessions alternating between Cue and No Cue conditions, both of which were initially unfamiliar, in the same fashion as the drug studies (ABABAB design). As these animals learned the response–reward contingency without the cue light, introduction of the cue was expected not to provide additional information about the reward; thus, according to theories of Pavlovian conditioning (see Dickinson, 1980), the cue should not have entered into association with the reward, and was therefore not predicted to affect choice.

Satiation. To exclude the interpretation that drug or delay effects were due to differences in primary motivation, the Houselight group were returned to their original signalling conditions, and were tested while alternating between hungry and sated states on consecutive days in the same manner as the drug/vehicle studies described above. Following a ‘hungry’ session, animals were placed on free food (lab chow) until the start of the next day’s ‘sated’ session, at which time the food was again removed for the ‘hungry’ session to follow. The comparison is therefore between animals on ~22 h food deprivation versus the sated state.

To establish whether prolonged deprivation had an effect on choice, a further satiation experiment was performed on the same subjects: half were placed on free food for a week while half remained hungry. They then performed the task for three sessions, after which the deprivation state was reversed for a week and a further three sessions’ data collected.

Descending delays. To demonstrate that the basic effect of delay did not depend on an ascending series of delays, the Houselight group were next trained under a descending series of delays (60, 40, 20, 10, 0 s) under their normal signalling conditions.

Omission of a cue. Following drug testing, the Cue group were tested with sessions that alternated between the Cue and No Cue conditions in an ABABAB design, and subsequently with an AAABBB design (three consecutive cue sessions followed or preceded by three no-cue sessions). The reason for this was as follows: It was expected that manipulations where the subjects were required to learn through their experience of the delays during the session (that is, manipulations that affected choice *retrospectively*) would be better detected by the AAABBB design, as this gives greater opportunity for expression of that learned behaviour under constant conditions. In contrast, this was not expected of manipulations that affected the subjects’ preference for delays that were about to occur (prospective choice; this distinction follows Killeen & Fetterman, 1988). While drugs are in principle capable of affecting choice prospectively, without requiring new learning, the only possible way that omission of the cue could affect choice behaviour is retrospectively: the subjects must learn that the cue no longer follows choice of the Delayed lever. In the ABABAB design, such learning might be obscured by the rapidly alternating contingencies.

Extinction. Following drug testing, the No Cue group were alternated between their normal task and extinction sessions, in which no reinforcement was delivered, in order to assess whether choice was controlled by a temporal stimulus (the passage of time within a session) or only by the exemplar (forced-choice) trials.

Statistical analysis

General statistical techniques were described in Chapter 2.

For baseline data, measures were calculated for each subject using pooled responses from all sessions, because an analysis using session as a within-subjects factor would reduce the power to detect effects of between-subjects factors (Bradley & Russell, 1998). Similarly, measures were calculated across the three session pairs of each drug study or behavioural manipulation. Choice ratios were calculated as the percentage of responses in which the Delayed lever was chosen, for free-choice trials only.

RESULTS

1. Acquisition and baseline performance

Acquisition of sensitivity to delay

In all groups, the rats' behaviour became sensitive to the delay following a number of training sessions (Figure 69A shows data for the Houselight group). In the first session, preference for the Delayed lever declined as the delays were introduced (not shown), presumably reflecting a degree of extinction as the delay was introduced. After this, preference for the delayed lever increased again until it was favoured at all delays. Finally, delay sensitivity was seen. It can be seen from Figure 69B that individual rats varied considerably in their preferences, despite the regular sampling of both levers at the start of each block.

Effect of cues on speed of acquisition

The presence of a cue during the delay speeded the acquisition of delay sensitivity. Following identical training procedures, the Houselight group reached criterion from session 11 (i.e. analysis of data from sessions 11–13, but not before, showed a significant effect of delay at $\alpha = .01$); the No Cue group met the criterion from session 18 and the Cue group from session 8. To confirm this effect statistically, the linear regression slopes (see *Methods*) for the first 14 sessions were subjected to an ANOVA. These slopes are shown in Figure 69C; analysis by group \times (session \times S) revealed a significant effect of session ($F_{8,149,138.531} = 6.021$, $\tilde{\epsilon} = .627$, $p < .001$), reflecting the acquisition of delay sensitivity, and a group \times session interaction ($F_{16,298,138.531} = 2.507$, $\tilde{\epsilon} = .627$, $p = .002$), indicating faster acquisition in the presence of a cue.

2. Baseline performance

Effect of cues on choice (between-subjects comparison)

All three groups reached a similar pattern of choice once they had satisfied the delay-sensitivity criterion (Figure 69D). There were no significant effects of the cue condition on choice (terms involving cue: $F_s < 1$, NS) though there was a significant effect of delay ($F_{1,905,40,002} = 38.489$, $\tilde{\epsilon} = .476$, $p < .001$). Similarly, there was no effect of cue on the regression slope measure (one-way ANOVA, $F < 1$, NS), even for the last baseline day ($F_{2,21} = 1.42$, NS). Taken on its own, this result suggests that the cue helps subjects to learn the contingencies in operation, but once these have been learned the cue plays no role in choice.

Omissions and latencies

Subjects' performance was reliable. Analysis across all groups showed that total omissions (failures to initiate a trial or respond on a lever) increased with delay ($F_{2,042,46.975} = 10.689$, $\tilde{\epsilon} = .511$, $p < .001$) and there was a significant but small tendency to slower initiation at long delays ($F_{2,283,47.937} = 8.632$, $\tilde{\epsilon} = .571$, $p < .001$), plausibly due to a degree of satiation. However, even at the final delay, omissions were only ~10%, or one out of the 12 trials (Table 19), despite the potential for rats to eat >10 g of pellets per session. Overall, of the 8400 choice trials analysed in the baseline data, subjects failed to initiate 5.7% of trials, and failed to respond to only 0.05% of initiated trials. Forced-choice trials were also responded to consistently: of the 1680 forced-choice presentations, 4.5% were not initiated and of those that were initiated, only 1.7% were not responded to.

Table 19: Omissions at different delays.

Delay (s)	% omissions (all kinds), mean \pm SEM
0	1.54 \pm 0.49
10	2.53 \pm 0.58
20	5.46 \pm 1.76
40	8.68 \pm 2.14
60	10.81 \pm 2.54

Subjects responded faster on the lever producing the large reinforcer ($F_{1,15} = 17.829$, $p = .001$) but this was independent of the delay and group ($F_s \leq 1.203$, NS). Food was collected within 10 s of delivery on 99.9% of rewarded trials. The latency to collect food was not affected by the delay, the cue condition, or the subject's preceding choice ($F_s < 1.327$, NS).




Figure 69. Task acquisition. **A:** Group means at different time points for the Houelight group. **B:** Individual records for the Houelight group, sessions 12–14, together with the group mean (*thick line*). **C:** Acquisition in different cue conditions as assessed by a regression slope measure (see text). *SED*, standard error of the difference between means for the group \times session interaction. The *SED* is the appropriate index of variation for comparison between different mean values (see e.g. Howell, 1997 for derivation). **D:** Responding under different cue conditions immediately prior to drug testing.

3. Pharmacological manipulations

In all drug studies, choice was analysed using an ANOVA with the model (dose \times delay \times S), and the main effect of delay on choice remained highly significant throughout ($p < .003$). While there appeared to be a small tendency for the within-session shift in preference to be more pronounced with prolonged experience of the task, there were no between-group differences in responding under vehicle for any drug/dose study (choice ratios, all F s < 1 ; slope measures, maximum $F_{2,21} = 2.42$, NS); thus, drug effects at each dose can be interpreted relative to the same group baseline. The use of a within-subjects design allows small drug effects to be detected, but the individual variability discussed above allows a strong interpretation — for a drug effect to be found, that drug must have consistent effects despite subjects' starting from different individual baselines.

Effects of d-amphetamine

Choice. The effects of amphetamine depended on the cue condition (Figure 70; Figure 71). In the Houselight group, amphetamine did not affect choice at any dose (main effects, $F_s < 1$; interactions with delay, $F_s < 2.08$, NS). In the No Cue group, amphetamine *reduced* preference for the large reinforcer at 1.0 mg/kg (drug \times delay interaction, $F_{4,28} = 3.336$, $p = .024$) and at 1.6 mg/kg (main effect of drug, $F_{1,7} = 6.834$, $p = .035$), but had no effect at 0.3 mg/kg (maximum $F_{1,7} = 3.30$, NS). In the Cue group, amphetamine *increased* preference for the large reinforcer at 0.3 mg/kg (main effect, $F_{1,7} = 12.393$, $p = .01$), and had no effect at other doses ($F_s < 2.25$, NS). The increase in preference for the large reinforcer caused by this dose, calculated as an arithmetical difference between choice ratios in the drugged and vehicle conditions, was 8.4% when averaged over all delays (ranging from a 2% increase at 20 s delay to an increase of 17.3% at 10 s). The only dose that produced a significantly delay-dependent effect was 1.0 mg/kg in the No Cue group, which significantly reduced choice ratios at 40 s delay ($p = .018$ by one-way ANOVA) but not at other delays ($p = .088$ at 20 s and $p > .266$ otherwise).

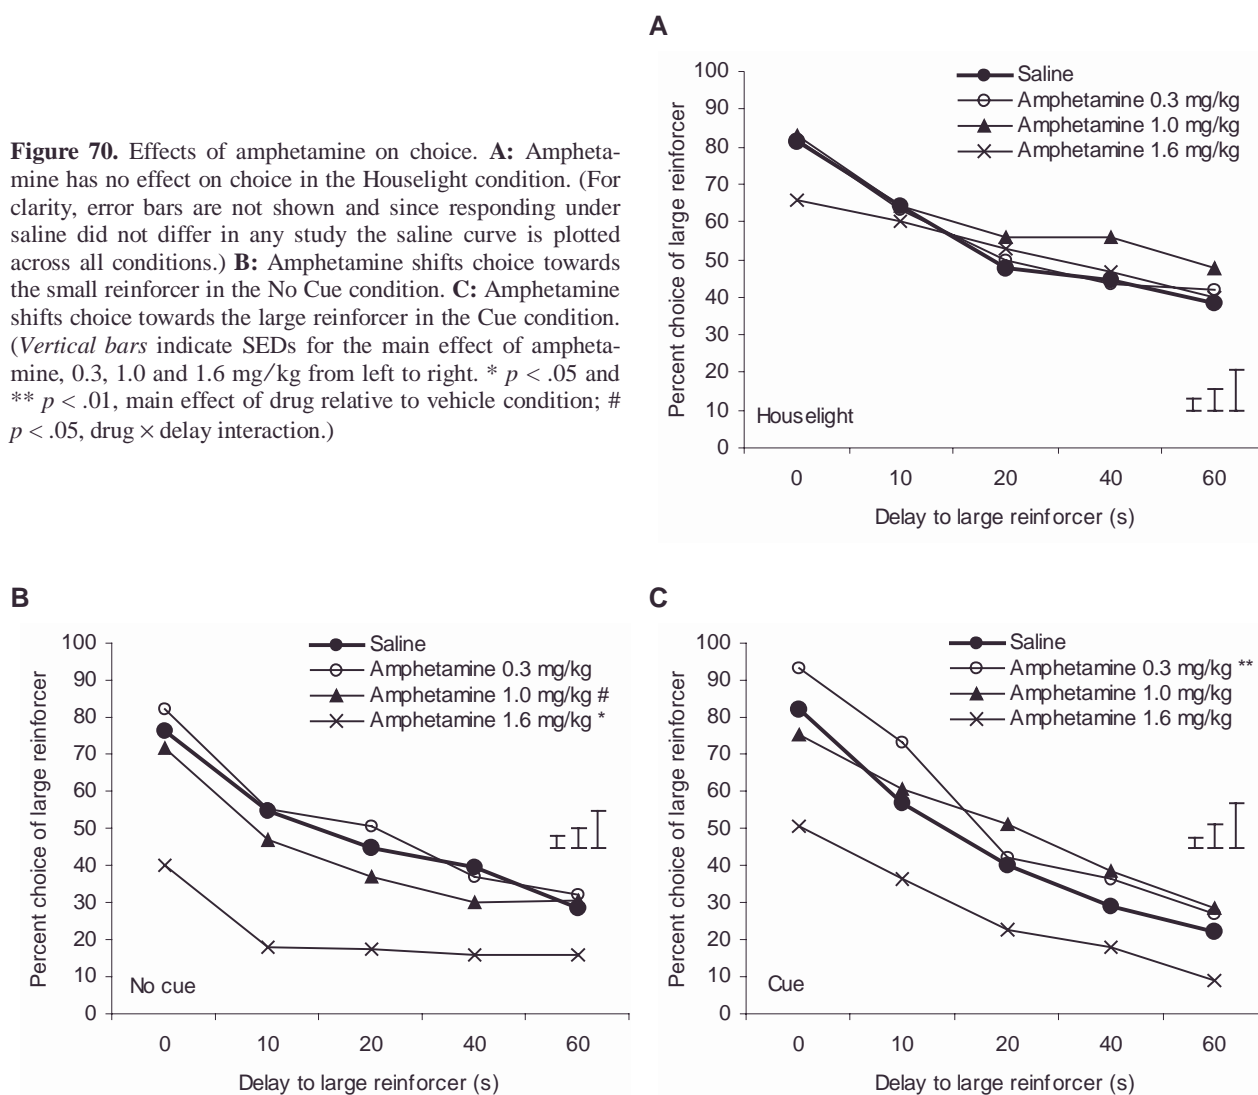
The effect of amphetamine to increase preference in the Cue group was not due to altered responding in the zero-delay condition. Firstly, although the absence of a drug \times delay interaction for 0.3 mg/kg strictly does not justify simple effects analyses, which also have lower power, such analyses showed that the effects at 10 s and 60 s (but not at 0 s) were significant in their own right. Secondly, elimination of the zero-delay condition from analysis did not alter the conclusion that 0.3 mg/kg caused a significant increase in choice ratios ($F_{1,7} = 9.801$, $p = .017$); the mean within-subject increase was 9.2% (as an arithmetical difference of %choice) in this analysis. Nonparametrically, six of eight rats showed an increase in preference for the delayed reinforcer calculated over all non-zero delays (Wilcoxon matched-pairs signed-ranks test, $p = .036$), and seven of eight rats showed an increase at the largest delay ($p = .025$). Nor did elimination of the zero-delay condition alter the conclusions about other doses.

The existence of a cue-dependent effect of amphetamine was confirmed statistically by testing data from the Cue and No Cue groups for a significant cue \times drug or cue \times drug \times delay interaction; this was found for 0.3 and 1.0 mg/kg (cue \times drug \times delay, $F_{8,112} = 2.498$, $p = .016$). The highest dose, 1.6 mg/kg, had marked effects on omissions and consequently did not demonstrate such an interaction. The functional relevance of the cue-dependent effect was assessed directly from the total mass of food obtained on choice trials at non-zero delays: 1.0 mg/kg amphetamine reduced the amount of food obtained by animals in the No Cue group by 10.7% (mean within-subject change from saline), but this dose caused the Cue group to obtain 12.7% *more* food.

Omissions. Only the highest dose of amphetamine increased omissions. As there were few omissions, the percentage of trials on which an omission (of the initiation or choice type) occurred was calculated and analysed independently of the delay. There was a significant overall effect of dose ($F_{1,286,27.008} = 24.709$, $\tilde{\epsilon} = .429$, $p < .001$), but no effect of cue (cue: $F_{2,21} = 2.465$, $p = .109$; cue \times dose: $F_{2,572,27.008} = 2.401$, $\tilde{\epsilon} = .429$, $p = .098$). Over all groups, the percentages of trials on which an omission occurred were 1.8 ± 0.3 (saline), 1.2 ± 0.6 (0.3 mg/kg), 1.9 ± 0.6 (1.0 mg/kg) and 15.7 ± 2.9 (1.6 mg/kg). Pairwise comparisons established that the 1.6 mg/kg dose differed from all other doses, which did not differ from each other.

Initiation latencies. Amphetamine slightly reduced initiation latencies at 0.3 mg/kg, and progressively increased them at higher doses. The mean initiation latencies in seconds (across all delays) were 1.207 ± 0.056 (saline), 1.057 ± 0.051 (0.3 mg/kg), 1.517 ± 0.108 (1.0 mg/kg), and 2.042 ± 0.139 (1.6 mg/kg). An analysis of data from all three groups revealed an effect of drug ($F_{1,969,45.289} = 32.905$, $\tilde{\epsilon} = .656$, $p < .001$), of delay ($F_{2,967,68.248} = 2.895$, $\tilde{\epsilon} = .742$, $p = .042$) and an interaction ($F_{4,513,103.805} = 4.38$,

Figure 70. Effects of amphetamine on choice. **A:** Amphetamine has no effect on choice in the Houselight condition. (For clarity, error bars are not shown and since responding under saline did not differ in any study the saline curve is plotted across all conditions.) **B:** Amphetamine shifts choice towards the small reinforcer in the No Cue condition. **C:** Amphetamine shifts choice towards the large reinforcer in the Cue condition. (Vertical bars indicate SEDs for the main effect of amphetamine, 0.3, 1.0 and 1.6 mg/kg from left to right. * $p < .05$ and ** $p < .01$, main effect of drug relative to vehicle condition; # $p < .05$, drug \times delay interaction.)



$\tilde{\epsilon} = .376$, $p = .002$), though this interaction was attributable to the fact that 1.6 mg/kg had a greater effect early on in the session (other doses had effects independent of the delay: an analysis without the highest dose showed no such interaction; $F < 1$, NS). Pairwise comparisons of the main effect of drug with a Sidak correction showed that all doses differed from each other ($p \leq .001$).

Choice latencies. The two higher doses (1.0 and 1.6 mg/kg) increased choice latencies, especially early in the session. An analysis across the three groups using the design (drug \times response \times delay \times S) revealed a significant drug \times delay interaction ($F_{3,644,21,986} = 4.833$, $\tilde{\epsilon} = .305$, $p = .007$). However, the effects of amphetamine did not depend on the response being made (drug \times response: $F_{3,18} = 2.767$, $p = .072$).

Nosepoking during the delay. Amphetamine dose-dependently reduced the proportion of the delay spent nose-poking in the food alcove from 16% (saline, mean across all delays) to 8% (1.6 mg/kg) ($F_{3,18} = 12.062$, $p < .001$; nose-poking data were unavailable for the Houselight group). In addition, independently of the effects of amphetamine, the presence of the cue supported higher levels of nose-poking, particularly at long delays (cue \times delay, $F_{3,18} = 4.519$, $p = .016$); the maximum effect occurred at 60 s delay, when the Cue group nose-poked for 16% of the delay (mean across all doses) and the No Cue group for 12%. This indicates that the cue had behavioural effects even in trained animals.




Figure 71. Effects of amphetamine on choice, replotted as mean within-subject changes for the Cue and No Cue groups. For each subject, the choice ratio under vehicle (as %choice of the large reinforcer) was subtracted from that under amphetamine, at each delay; the ordinate is the mean of these values. The effects of amphetamine on the Cue and No Cue groups differed significantly at 0.3 and 1.0 mg/kg (see text).

Food collection latencies. In this study, subjects collected the immediate reward faster than the delayed reward; neither amphetamine nor the delay had any influence on collection latency. An analysis across all groups using the model (drug \times response \times delay \times S) revealed a main effect of response (response: $F_{1,6} = 10.53$, $p = .018$), but no other terms were significant (maximum $F_{1,139,6.835} = 3.351$, $\tilde{\epsilon} = .38$, $p = .109$).

To summarize, at doses that did not grossly alter responding, the presence of a cue altered the effects of amphetamine on choice. Amphetamine had a cue-independent effect to reduce preference for the delayed reinforcer, and a cue-dependent effect to increase preference.

Effects of chlordiazepoxide

Choice. Chlordiazepoxide (CDP) generally promoted choice of the Immediate lever, and its effects did not alter in the presence of a cue (Figure 72). CDP had effects at all doses used except 1.0 mg/kg. Half a session's worth of data from one subject in the No Cue were lost from the 10 mg/kg study due to a malfunction.

In the No Cue group, chlordiazepoxide promoted choice of the smaller reinforcer, but only at 10 mg/kg ($F_{1,7} = 14.876$, $p = .006$), a dose that also increased the omission rate (see below); it had no effect on choice at other doses (closest to significance: main effect for 3.2 mg/kg, $F_{1,7} = 3.424$, $p = .107$). In the Cue group, the effects varied according to the dose of CDP and the delay. At 10 mg/kg the effect was similar to that for the No Cue group but not significant ($F_{1,6} = 5.729$, $p = .054$). However, 5.6 mg/kg caused a smaller but highly significant shift towards the small reinforcer (drug \times delay interaction, $F_{4,28} = 2.871$, $p = .041$; main effect of drug: $F_{1,7} = 17.414$, $p = .004$), an effect that was significant at 10- to 40-s delays (simple effects, $p \leq .024$) but not for 0 or 60 s ($p \geq .058$). At 3.2 mg/kg, CDP had mixed effects (drug \times delay interaction, $F_{4,28} = 2.843$, $p = .043$), promoting choice of the large reinforcer at 10 s ($F_{1,7} = 6.973$, $p = .033$) and of the small reinforcer at 40 s ($F_{1,7} = 6.831$, $p = .035$); effects at other delays were not significant ($p \geq .07$). The lowest dose, 1.0 mg/kg, had no effect in either group (maximum $F_{1,7} = 2.956$, $p = .129$).

Overall, no evidence for a cue-dependent effect of CDP was found. As before, data from the Cue and No Cue groups were tested for a cue \times drug or cue \times drug \times delay interaction: no such terms were significant.

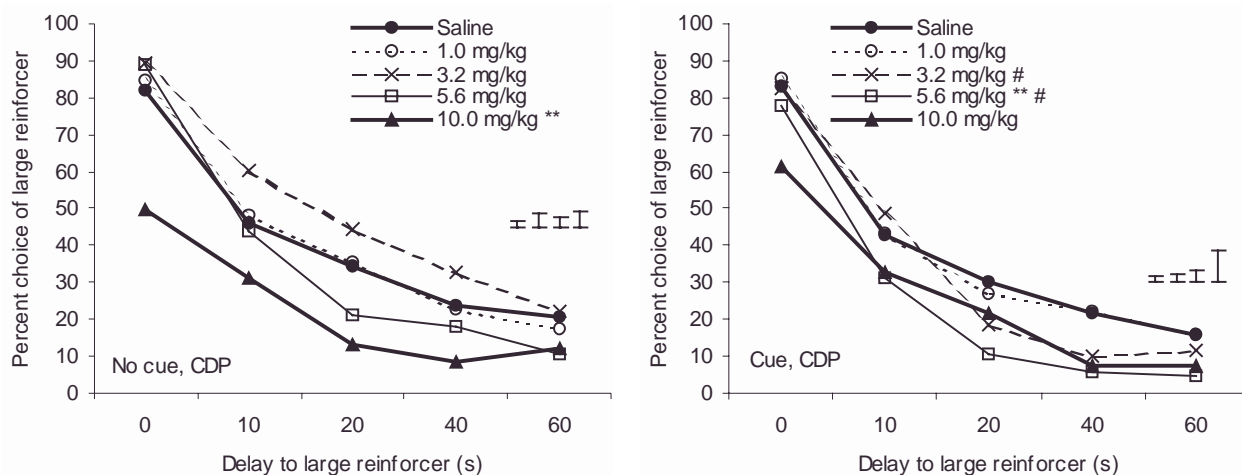


Figure 72. Effects of chlordiazepoxide on choice. As before, each line represents the mean of 8 subjects' choice ratios, calculated for three drugged sessions each, except the saline curve, which is calculated across all four dose studies (12 sessions) for simplicity of presentation as responding under saline did not differ in any dose study or between signalling conditions. (Vertical bars indicate SEDs for the main effect of chlordiazepoxide, 1.0, 3.2, 5.6 and 10.0 mg/kg from left to right. ** $p < .01$, main effect of drug relative to vehicle condition; # $p < .05$, drug \times delay interaction.)

Omissions. Only the highest dose (10 mg/kg) markedly increased omissions ($F_{1,187,16,611} = 24.442$, $\tilde{\epsilon} = .297$, $p < .001$). Indeed, this dose induced obvious somnolence in a number of subjects within minutes of administration. The percentages of trials on which an omission occurred were 1.4 ± 0.3 (saline), 1.2 ± 0.3 (1.0 mg/kg), 1.6 ± 0.4 (3.2 mg/kg), 4.1 ± 1.3 (5.6 mg/kg) and 33.2 ± 6.0 (10.0 mg/kg). Pairwise comparisons showed that 10.0 mg/kg differed from all other doses ($p \leq .004$ in all cases) but no other doses differed from each other ($p \geq .255$).

Initiation latencies. The highest dose (10 mg/kg) increased initiation latencies, particularly at the start of the session, but no other dose had an effect. Mean initiation latencies in seconds were 0.994 ± 0.093 (saline), 1.02 ± 0.088 (1.0 mg/kg), 0.883 ± 0.072 (3.2 mg/kg), 0.958 ± 0.06 (5.6 mg/kg), and 1.823 ± 0.149 (10 mg/kg). An analysis by cue \times (drug \times delay \times S) revealed a main effect of drug ($F_{1,61,20,926} = 29.785$, $\tilde{\epsilon} = .401$, $p < .001$) and a drug \times delay interaction ($F_{7,803,101,442} = 3.12$, $\tilde{\epsilon} = .488$, $p = .004$); no other terms were significant (drug \times delay \times cue: $F_{7,803,101,442} = 1.934$, $\tilde{\epsilon} = .488$, $p = .064$; other terms: $F < 1$, NS). Pairwise comparisons showed that 10 mg/kg differed from all other doses ($p \leq .001$), which did not generally differ from each other (1 versus 3.2 mg/kg, $p = .049$; all other comparisons, $p \geq .171$). An analysis without the data for 10 mg/kg did not exhibit any delay-dependent drug effects ($F_s < 1.118$, NS).

Choice latencies. The pattern of results was identical to that for initiation latencies. There were insufficient data to analyse using a full model with cue and response as factors, so (drug \times delay \times S) was used. This revealed a main effect of drug ($F_{1,752,24,533} = 82.666$, $\tilde{\epsilon} = .438$, $p < .001$) and a drug \times delay interaction ($F_{3,213,44,978} = 11.472$, $\tilde{\epsilon} = .201$, $p < .001$) but no effect of delay ($F < 1$, NS). Pairwise comparisons showed that 10.0 mg/kg differed from all other doses ($p < .001$), which did not differ from each other ($p \geq .068$).

Nosepokes during the delay. CDP did not have consistent effects on nosepoking. An ANOVA by cue \times (dose \times delay \times S) revealed a complex pattern of results, there being a dose \times delay \times cue interaction ($F_{4,852,19,407} = 4.621$, $\tilde{\epsilon} = .404$, $p = .006$). However, inspection of the data revealed that these results were entirely due to an aberrant increase in nosepoking at 40 s under 10 mg/kg in the Cue group; analysis without the 10 mg/kg data showed no significant effects of any term ($p > .093$).

Food collection latencies. CDP did not affect the latency to collect the reward. Again, there was insufficient data to use a full model, so (drug \times delay \times S) was used; no terms were significant (drug: $F_{1,264,17,693} = 1.636$, $\tilde{\epsilon} = .316$, NS; delay: $F_{2,205,30,867} = 2.401$, $\tilde{\epsilon} = .551$, NS; drug \times delay: $F_{1,7,23,795} = 1.263$, $\tilde{\epsilon} = .106$, NS).

Effects of α -flupenthixol

Choice. α -Flupenthixol had a weak effect to promote choice of the small reinforcer, irrespective of the cue condition (Figure 73). This effect reached significance for the No Cue group at 0.125 mg/kg (main effect, $F_{1,7} = 6.805$, $p = .035$) and for the Cue group at 0.25 mg/kg ($F_{1,7} = 8.204$, $p = .024$); though this effect was statistically independent of delay, it was numerically greatest at delays of 20–60 s. No other effects were significant, though there was a tendency for 0.125 mg/kg to promote choice of the small reinforcer in the Cue group as well ($F_{1,7} = 4.415$, $p = .074$). The pattern of choice remained remarkably stable at high doses despite a large increase in omissions (see below).

α -Flupenthixol had a greater effect to decrease choice ratios in the Cue condition than in the No Cue condition at 0.125 mg/kg: in addition to a main effect of α -flupenthixol to decrease choice ratio scores ($F_{1,14} = 7.846$, $p = .014$), there was a cue \times drug \times delay interaction ($F_{4,56} = 2.671$, $p = .041$). Analysis of simple effects of drug at different delays showed that this interaction was due to a greater effect of 0.125 mg/kg to decrease choice ratios in the Cue than in the No Cue group at 40 s delay (simple cue \times drug interaction, $F_{1,14} = 7.597$, $p = .015$). However, this cue-dependent effect was small and there were no such effects at 0.25 and 0.5 mg/kg.

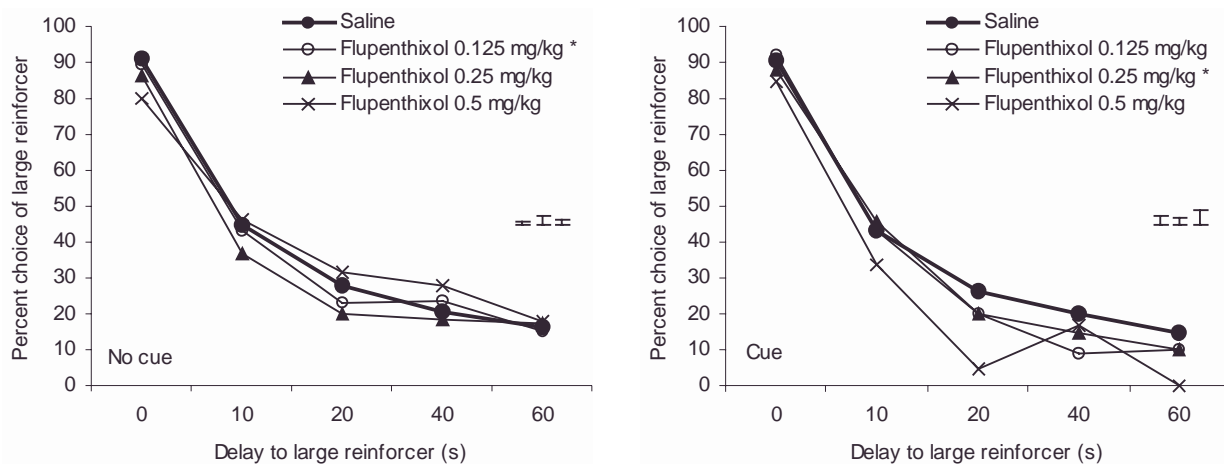


Figure 73. Effect of α -flupenthixol on choice. As before, each line represents the mean of 8 subjects' choice ratios, calculated for three drugged sessions each, except the saline curve, which is calculated across all four dose studies for simplicity of presentation as responding under saline did not differ in any dose study or between signalling conditions. (Vertical bars indicate SEDs for the main effect of α -flupenthixol, 0.125, 0.25 and 0.5 mg/kg from left to right. * $p < .05$, main effect of drug relative to vehicle condition.)

Omissions. The higher doses of α -flupenthixol increased omissions ($F_{1,405,19,676} = 73.813$, $\tilde{\epsilon} = .468$, $p < .001$); this was independent of the cue ($F_s < 1$). The percentages of trials on which an omission occurred

were 2.3 ± 0.6 (saline), 3.0 ± 1.1 (0.125 mg/kg), 7.3 ± 1.8 (0.25 mg/kg) and 44.4 ± 4.6 (0.5 mg/kg). Pairwise comparisons showed that 0.5 mg/kg differed from all other doses ($p < .001$ in all cases); in addition, 0.25 mg/kg differed from saline ($p = .026$) but no other doses differed from each other ($p \geq .145$).

Initiation latencies. Flupenthixol dose-dependently increased initiation latencies, particularly at long delays, late in the session. An ANOVA by cue \times (drug \times delay \times S) revealed main effects of drug ($F_{1.835,14.683} = 4.678$, $\tilde{\epsilon} = .612$, $p = .029$) and delay ($F_{1.746,13.969} = 4.065$, $\tilde{\epsilon} = .437$, $p = .045$) and a drug \times delay interaction that escaped significance ($F_{2.664,21.308} = 3.091$, $\tilde{\epsilon} = .222$, $p = .054$). No other terms were significant ($F_s \leq 1.221$, NS).

Choice latencies. Flupenthixol's effects on choice latencies were similar to those on initiation latencies, with an increase in latency particularly at long delays. As there were insufficient data for a full model at all doses, the design cue \times (drug \times delay \times S) was used. There were main effects of drug ($F_{2.405,19.242} = 7.854$, $\tilde{\epsilon} = .802$, $p = .002$) and delay ($F_{4,32} = 11.205$, $p < .001$) and a significant drug \times delay interaction ($F_{5.118,40.945} = 8.098$, $\tilde{\epsilon} = .427$, $p < .001$). In addition, there was a cue \times drug \times delay interaction ($F_{5.118,40.945} = 2.486$, $\tilde{\epsilon} = .427$, $p = .046$). No other terms were significant ($p \geq .094$). However, simple interaction analyses did not reveal a dose whose effects were demonstrably different in the Cue and No Cue groups (delay \times cue interactions, $p > .128$).

Nosepokes during the delay. α -Flupenthixol dose-dependently blocked the ability of the cue to sustain higher rates of nose-poking (Figure 74). In this analysis, the number of omissions at 0.5 mg/kg was so high that it was necessary to omit these data for analysis of the other doses. This revealed a dose \times cue interaction ($F_{1.343,10.748} = 9.573$, $\tilde{\epsilon} = .672$, $p = .007$) in addition to main effects of dose ($F_{1.343,10.748} = 19.636$, $\tilde{\epsilon} = .672$, $p = .001$) and cue ($F_{1,8} = 9.465$, $p = .015$), and a dose \times delay interaction ($F_{6,48} = 5.645$, $p < .001$); no other terms were significant ($F_s < 1.4$, NS). Analysis of simple effects of the drug (across all delays) showed that in the No Cue group, subjects' nose-poking was unaffected by flupenthixol, while nose-poking was significantly reduced by the 0.25 mg/kg dose in the Cue group.




Figure 74. Effect of α -flupenthixol on nose-poking during the delay to reinforcement. The **left** panel shows data from both groups (SED, standard error of the difference for the main effect of α -flupenthixol; ** $p < .01$, difference from saline). The **right** panel, plotting data averaged across all delays, illustrates that the normal ability of the cue to sustain nose-poking was abolished by α -flupenthixol (SED, standard error of the difference for the dose \times cue interaction; ## $p < .01$ for this interaction).

Food collection latencies. α -Flupenthixol did not affect the latency to collect reward. Data were analysed using the model cue \times (drug \times delay \times S); aside from slightly longer collection latencies at long delays (main effect of delay: $F_{2.258,18.036} = 3.677$, $\tilde{\epsilon} = .564$, $p = .041$), no terms were significant ($p \geq .141$).

4. Behavioural manipulations

Omission of delays

Omission of delays had clear effects to increase preference for the large reinforcer (Figure 75; Houselight group). There were significant effects of the Delay/No Delay factor ($F_{1,7} = 7.802, p = .027$), trial block ($F_{2,023,14,159} = 17.005, \tilde{\epsilon} = .506, p < .001$) and a significant interaction ($F_{1,589,11,121} = 8.094, \tilde{\epsilon} = .397, p = .009$). The effect of omitting the delays was not complete, as subjects still altered their preference across the session in the absence of any delays (simple effect of trial block in the No Delay condition, $F_{4,28} = 6.736, p = .001$).




Figure 75. Effect of removing all delays on choice (Houselight group). (*SED*, standard error of the difference for the interaction term; ** $p < .01$ for this term.)

The development of a preference for the large reinforcer throughout a no-delay session was not immediate. Figure 76 shows the manner in which preferences changed across individual sessions (note that the data set is slightly different from that in Figure 75). Inspection of this figure suggests that the typical within-session shift towards the small reinforcer was present for the first no-delay session and was gradually eliminated. Indeed, analysis of the data from the three no-delay sessions using the model (session pair₃ × trial block₅ × S) demonstrated a change in the pattern of responding across the sessions ($p = .013$), with subjects exhibiting a within-session shift in preference on the first session of the three ($p = .006$) but not the third ($p = .241$). Subjects always exhibited always exhibited a within-session preference shift within delay sessions ($p \leq .003$), although it can be seen that the pattern of responding changed in these sessions too ($p < .001$). In particular, the alteration in the change in preference at from 0 s to 10 s delay during Delay sessions is interesting, as the 10-s block is first block during which the rat can determine whether it is experiencing a Delay or a No Delay session. If the rat were not successfully using the forced-choice trials as exemplars for the rest of the block, one would expect such ‘contamination’ of performance in delay sessions by experience of sessions without delays.




Figure 76. Group means for Houselight group ($n = 8$) when delays were omitted. The data are those shown in Figure 75, except that the session preceding the previous test is included, and the last session of the test omitted, for those rats who began the previous test with the No Delay condition. ($2SED$, twice the standard error of the difference for the session pair \times delay/no-delay \times trial block interaction; see text.)

Effect of cues on choice (within-subjects comparison)

While Figure 69D (p. 178) showed the effects of the cue condition on choice in a between-subjects comparison, a more sensitive test is a within-subjects comparison; not only does this have increased statistical power, but it reduces the potential for a learned adaptation to compensate for underlying cue effects on choice.

Introduction of a cue

The Houselight group were trained with successive sessions that alternated between Cue and No Cue conditions (both of which were initially unfamiliar) in the same fashion as the drug studies. As predicted, the cue had no effect on choice, even when the manipulation was extended to twelve sessions (Figure 77); analysis showed $F < 1$ (NS) for all terms involving cue.




Figure 77. Lack of effect of introducing a cue for the Houselight group. The first six such test sessions are shown; no differences were observed in the second six sessions. (SED , standard error of the difference for the interaction term.)

Omission of a cue

Removing the cue from the Cue group reduced subjects' ability to choose the large, delayed reinforcer, although only when subjects experienced several consecutive sessions without the cue. Omitting the cue in alternate sessions (ABABAB design) did not affect choice ($F \leq 1.255$, NS for all terms involving cue).

However, when the Cue group experienced three consecutive cue sessions followed or preceded by three no-cue sessions (AAABBB design), an effect of cue emerged. The cue supported more frequent choice of the large reinforcer, particularly at long delays (Figure 78). An analysis of choice ratios as (cue \times delay \times S) showed a significant cue \times delay interaction ($F_{2.636,18.451} = 3.564$, $\tilde{\epsilon} = .659$, $p = .039$). Examination of individual subjects' performance showed that at every non-zero delay, six out of eight rats showed more frequent choice of the large reinforcer in the presence of the cue.




Figure 78. Effect of removing a cue from the Cue group: three consecutive cue sessions were given, followed by three consecutive no-cue sessions (or the reverse order, counter-balanced). (*SED*, standard error of the difference for the delay \times cue/no cue interaction; * $p < .05$ for this term.)

Effects of prefeeding

Sating the subjects by giving 22 h free access to food had no effect on choice, despite progressively increasing initiation latencies through the session (Figure 79). Analysis of choice using the model (hunger \times delay \times S) showed no significant terms involving hunger ($F < 1$, NS). Every animal made more omissions when sated (heterogeneity of variance necessitated a non-parametric test: Wilcoxon matched-pairs signed-ranks test, $p = .012$). Initiation latencies were reliably increased by satiation: an analysis of variance using the model (hunger \times delay \times S) revealed a main effect of hunger ($F_{1,7} = 12.368$, $p = .01$) and hunger \times delay ($F_{1,969,13,781} = 5.269$, $\tilde{\epsilon} = .492$, $p = .02$), with no main effect of delay ($F_{1,474,10,321} = 2.378$, $\tilde{\epsilon} = .369$, NS).

Prolonged satiation or deprivation had no effect on choice (Figure 79D). Maintenance on a more severe food deprivation regimen for a week reduced body mass to 86.1% of that following a week's free access to food (mean within-subject change), yet the effect of deprivation on choice was not significant ($F_s \leq 1.38$, NS).

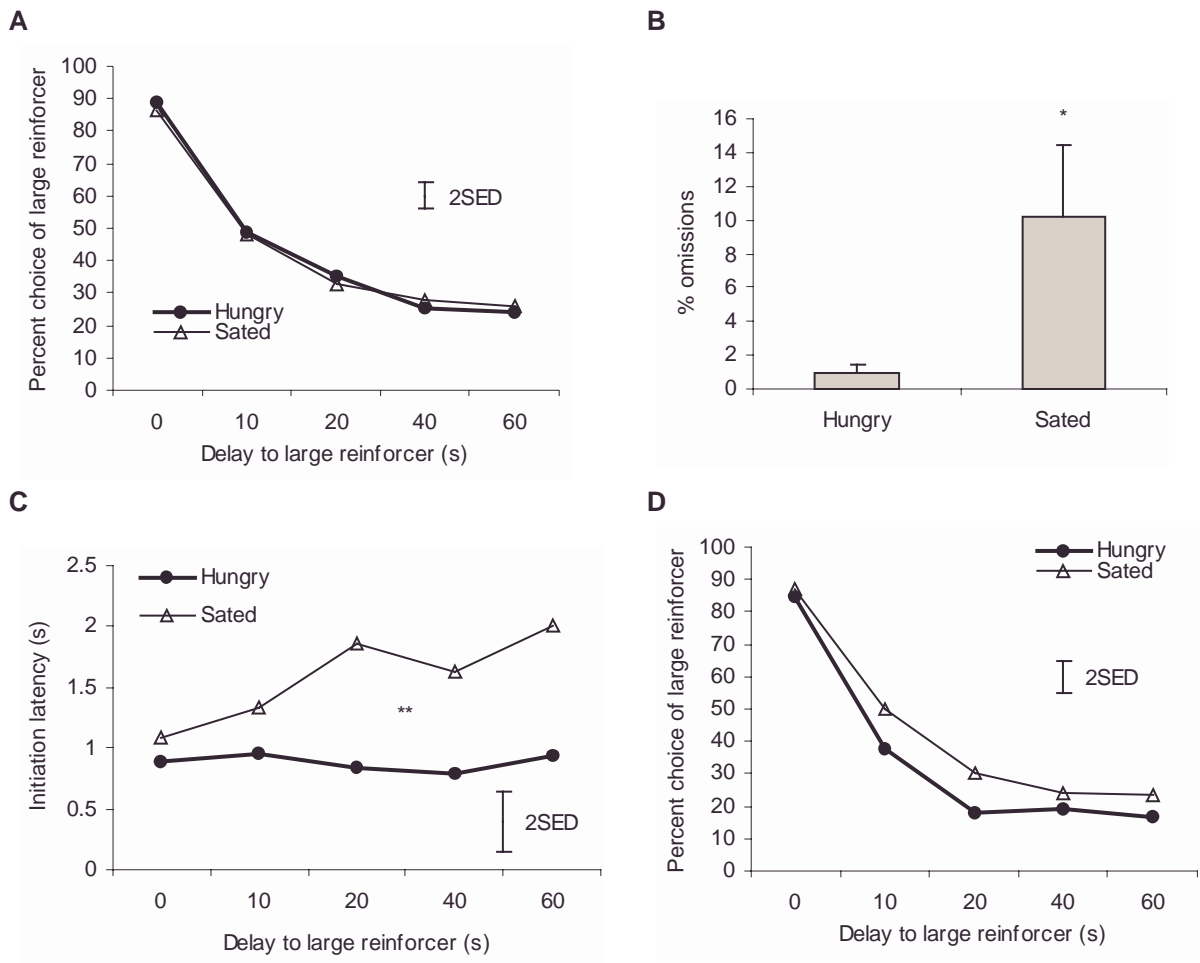


Figure 79. Effects of prefeeding. **A:** Satiation had no effect on choice. **B:** Satiation increased the number of omissions made. **C:** Satiation increased initiation latencies as the session progressed. **D:** Choice was not affected even when subjects were maintained in each deprivation condition for a week prior to testing. (*2SED*, twice the standard error of the difference for the hunger \times delay term; * $p \leq .05$, ** $p \leq .01$ for the effect of hunger.)

Descending delays

Changing from an ascending to a descending series of delays reversed the direction of the subjects' preference shift within the session (Figure 80); the preference shift does not therefore depend on the use of an ascending series of delays. After the change, the group took 11 sessions to re-satisfy the stability criterion, suggesting that trained animals adjust their responding to a new pattern of delays at a similar speed to naïve subjects.




Figure 80. Reversal of delays. **Left panel** shows the effect of delay reversal on choice (Houselight group). Each line represents data from three consecutive sessions. The curve labelled ‘before reversal’ was part of the long-term satiation experiment, and met the stability criterion (effect of delay significant, $p < .01$). The first three sessions also met this criterion, but in the (now) inappropriate direction. After this, no set of sessions met the stability criterion until post-reversal sessions 11–13, and subsequently. **Right panel** shows mean regression slopes calculated for each session individually (see Methods). The reversal renders the subjects’ preference ‘incorrect’ (shift from appropriate negative slope to inappropriate positive slope) and this slope declines gradually back towards the previous level.

Extinction

Extinction increased the number of omissions (from 4.8 ± 2.7 to 33.0 ± 8.3 per session; $F_{1,7} = 16.7, p = .005$). Extinction also affected choice in that preference tended towards indifference (50% ratio; Figure 81). However, an effect of delay remained in extinction: preference for the large reinforcer still declined throughout the session. Thus, extinction caused the animals to respond infrequently and randomly, but their tendency to choose the lever formerly associated with large reinforcement persisted for the first block despite the forced-choice trials preceding it. An analysis of choice ratios by (extinction \times delay \times S) showed effects of extinction ($F_{1,7} = 6.83, p = .035$), delay ($F_{4,28} = 36.5, p < .001$) and extinction \times delay ($F_{4,28} = 6.98, p < .001$). There was also a simple effect of delay in the Extinction condition ($F_{3,18} = 7.20, p = .003$), in which responding differed significantly from 50% choice in the first block (one-sample t test: $t_7 = 5.13, p = .001$) but during no other block ($|t| < 1.01, NS$).




Figure 81. Choice in extinction. (SED, standard error of the difference for the interaction term; *** $p < .001$ for this term.)

DISCUSSION

The effects of amphetamine on impulsive choice depended strikingly upon whether the delayed reward was signalled, with amphetamine increasing impulsivity in the unsignalled condition and decreasing impulsivity when the delay was bridged by a signal. The dopamine receptor antagonist α -flupenthixol had opposite, although less marked effects in the cued condition than amphetamine. In contrast, effects of chlordiazepoxide on choice were not influenced by signalling the delayed reward. In order to interpret these results, the factors controlling baseline performance will first be considered.

Task validation

This work replicates and extends the findings of Evenden and Ryan (1996) concerning performance on this task. Subjects' choice behaviour gradually came under the control of the programmed delay during training, even though the overall rate of reinforcement on each lever never varied and the optimal strategy was always to choose the Delayed lever. Rats remained sensitive to the delays even after prolonged training. The within-session shift in preference was not due to satiation or fatigue: not only did animals reliably collect food even at the end of the session, but prefeeding and prolonged changes in deprivation state failed to affect choice behaviour (in agreement with Richards (1997b), though not with Bradshaw and Szabadi (1992) or Ho *et al.* (1997)). Removing the delays dramatically increased rats' preference for the large reinforcer, compared to the same time point in a normal session. Finally, when subjects were switched from an ascending to a descending series of delays, their preference came to shift in the opposite direction.

Some rats were far from 100% choice of the large reinforcer at zero delay; this differs from typical findings with discrete-trial and ratio schedules, where maximization is the norm (see Mackintosh, 1974, pp. 190–195). The departure from 100% was also greater than that found by Evenden and Ryan (1996); these authors always exposed rats to the differences in reinforcer magnitude before delays were introduced, whereas in the present study both were introduced simultaneously. The training procedure of Evenden & Ryan (1996) allows better establishment of the contingency between the 'large' lever and its reward, given that acquisition of this contingency may be impaired by the delay (Dickinson *et al.*, 1992), but establishes a bias for the large lever by the time the delays are introduced. In addition, Evenden and Ryan (1996) used a greater difference in reinforcer magnitudes between the two levers.

Some subjects were initially biased towards the Immediate lever and some towards the Delayed lever. Casual observation suggested that subjects tended to be biased towards the right-hand lever in each case, which was the lever trained second in the initial FR schedule. The training procedure was for a criterion number of responses in a fixed *time*, so subjects accumulated more responses on the second lever trained as the instrumental contingency was already established and generalized from the first lever. This probably encourages the development of habitual responding on the second lever to be trained; such habits depend on the number of reinforcers obtained (Dickinson *et al.*, 1995), and the effect might be avoidable by training to a criterion total number of reinforcers per lever. It is of some interest that Tomie *et al.* (1998) used exactly the same training procedure as the present study, and effectively the No Cue condition (the group slowest to acquire delay sensitivity in this experiment), and found a number of rats with extreme positional bias ('delay-insensitive'). However, this effect can only decrease the power of tests to find effects of neuropharmacological manipulations on choice.

The rats' persistence in shifting their responding from the Delayed to the Immediate lever during sessions when all delays were zero, and during extinction sessions, implies that they failed fully to use the forced-choice trials as exemplars for the subsequent block of choice trials (thus, performance on the task

cannot be accurately characterized as ‘fully-informed choice’). It suggests strongly that the passage of time or trials acted as a discriminative stimulus that came to control responding, because when all delays are zero, no other stimulus is likely to cause a shift in choice from large to small reward. Subjects may take many sessions to acquire the characteristic within-session shift in choice, and to reacquire criterion performance when the delay sequence is reversed, in part because they must learn a session-wide temporal discrimination.

Role of signals present during the delay

The acquisition of delay sensitivity was facilitated by the presence of a discrete cue signalling the delayed reward, presumably by promoting discrimination between early and late trials and by speeding learning of the instrumental contingencies. This cue had no gross effects on stable choice performance, although it supported a higher rate of nose-poking in the food alcove during the delay. However, removing the cue demonstrated that it promoted or supported choice of the large reinforcer in animals that learned the task in its presence, despite extensive experience with the task (these subjects were nearing their hundredth session) and with no differences in primary reinforcement. The effects of cue omission were manifest only when animals had the opportunity to learn over several sessions that the cue was no longer contingent upon responding, as observed for other schedules controlled by response-contingent stimuli predicting reward (Everitt *et al.*, 1989; Arroyo *et al.*, 1998).

The effects of the stimulus light cue are notable, because there are several environmental stimuli that could provide information to the subjects about the impending food reward. The absence of the small reinforcer following choice is an unambiguous signal for all subjects that the large reward is imminent, but an even more obvious cue is having just responded on the Delayed lever (see also Garrud *et al.*, 1981). In the Houselight condition, the houselight was paired in an overlapping fashion with both the large and the small reward, and was also present at the start of the trial. Unambiguous interpretation of this group’s results is therefore difficult. However, the Cue and No Cue conditions differed in only one respect: the presence or absence of a stimulus light preceding the large, delayed reward.

The results were entirely consistent with the cue being a conditioned reinforcer. Confirmation of this would require demonstration that the effect of the stimulus was due to its association with primary reinforcement, that the effect on behaviour was a consequence of the response–stimulus contingency, and that the response had never produced primary reinforcement (Mackintosh, 1974, p. 234). Such a demonstration would also be required to be certain that the cue did not become aversive by virtue of the long delay to reward. However, the conditions existed for a positive Pavlovian association to form between the predictive cue and the large reinforcer: the reward was delivered at a range of times after the onset of the stimulus, and never in its absence. The absence of an effect of introducing a cue light to the Houselight group, for whom it provided no extra information, argues against a simple ‘stimulus-seeking’ explanation of the cue’s effects in undrugged animals. Faster acquisition of delay sensitivity in the presence of conditioned reinforcement is to be expected if such sensitivity is a consequence of discrimination learning (Grice, 1948), as suggested above. Finally, the absence of the cue at the moment of choice precludes its role as a discriminative stimulus in the usual sense; whether the cue acts as a conditioned reinforcer by acquiring some properties of the reinforcer or by providing information about its availability is a separate question (see Mackintosh, 1974, pp. 250–259).

Effects of *d*-amphetamine

As predicted, amphetamine had a dual effect on choice of delayed reinforcement, comprising a cue-independent effect to reduce preference for the large, delayed reward, and a cue-dependent effect to increase this preference. Relative to vehicle conditions, the magnitude of the amphetamine increase in the Cue condition was moderate, altering an average of 9% of subjects' choices from 'immediate' to 'delayed' reward decisions at non-zero delays. However, comparing the Cue and No Cue groups showed that the cue made a large difference to the effects of 1.0 mg/kg amphetamine on responding, altering an average of 16% of decisions from 'immediate' to 'delayed' choices. The cue-dependent effect of amphetamine to increase preference for the delayed reinforcer was consistent across subjects and resulted in a substantial increase in the amount of food earned by the Cue group.

These effects of amphetamine are consistent with previous work on impulsive choice, and may explain certain discrepancies in the literature: Evenden and Ryan (1996) used a task equivalent to the No Cue condition in the present study and found that amphetamine reduced preference for the large, delayed reward. The opposite result has been obtained using the adjusting-amount procedure (Richards *et al.*, 1997b), in which subjects make repeated choices between an immediate, variable amount of water and a delayed large reinforcer. Richards *et al.* sounded a tone for the duration of the delay, analogous to the Cue condition here, and have shown that amphetamine and the amphetamine analogue methamphetamine increase preference for the larger, delayed reward (Richards *et al.*, 1997a; 1999; Wade *et al.*, 2000). It is therefore clear that signals during the delay must be taken into account in future research on delayed reinforcement.

It is suggested that the cue-dependent effect of amphetamine reflects the potentiation of conditioned reinforcing properties of the cue, which predicts the arrival of a large reward. The efficacy of conditioned reinforcers is selectively increased by amphetamine and related compounds (Hill, 1970; Robbins, 1976; Robbins *et al.*, 1983) and this effect depends on a predictive relationship between the conditioned stimulus (CS) and the primary reinforcer (Robbins, 1976; Robbins & Koob, 1978). In the present study, the cue supported choice of the large reinforcer in animals trained in its presence, and amphetamine potentiated this effect; it is conceivable that the impulsivity-reducing effects of amphetamine in this task were entirely due to its actions to increase the efficacy of conditioned reinforcement. The neural locus for the impulsivity-reducing effects of amphetamine remains to be established, though this hypothesis predicts that it would be the nucleus accumbens shell as this is the critical site for the potentiation of conditioned reinforcement by amphetamine (Taylor & Robbins, 1984; Parkinson *et al.*, 1999b), a drug whose systemic effects in this respect are relatively weak (Robbins *et al.*, 1983). The finding that amphetamine's effects depended on the training history of the subjects is also analogous to that of Terrace (1963), who suggested that drug effects on S+/S- discrimination depended upon whether the training procedure established the S- as aversive; in the present study, the cue-dependent effects of amphetamine are hypothesized to depend on training that establishes the cue as an appetitive stimulus.

It is unlikely that this result simply represents another instance of the phenomenon that behaviour controlled by external stimuli is less susceptible to disruption by amphetamine (Laties & Weiss, 1966; Carey & Kritkowsky, 1972; Laties, 1972). Firstly, it should be noted that amphetamine might fail to disrupt behaviour controlled by external stimuli because it potentiates the effects of conditioned reinforcers, rather than because it improves discriminative stimulus control (Laties *et al.*, 1981), and there is little evidence to suggest that amphetamine facilitates control by purely discriminative (noncontingent) stimuli (e.g. Moerschbaecher *et al.*, 1979) or promotes responding for informative stimuli that are not themselves paired with reward (Branch, 1975). Secondly, the fact that amphetamine *increased* preference for the

large reinforcer in the presence of the cue implies that the cue does more than ameliorate an amphetamine-induced deficit.

One other interpretation deserves consideration. At the point when drug testing began, all groups had attained the same degree of control of behaviour by the delays. Nevertheless, as the cue affected the speed of task acquisition, the effects of each dose were assessed at different time points relative to the start of training in each group (earliest in the cued group). These slight temporal differences might thus account for the observed differences in the effects of amphetamine between the cued and uncued groups. However, this seems unlikely, as direct comparison of the vehicle data for each dose studied revealed no differences whatsoever in responding between the groups.

The cue-independent effect of amphetamine might reflect some specific psychological process. For example, amphetamine has been suggested to increase the speed of an 'internal clock' (Meck, 1983; Gibbon *et al.*, 1997); this might have affected choice prospectively (i.e. the subject perceives itself to be at a later time-point in the session than it actually is, hastening the within-session shift towards the Immediate lever), or it may have affected retrospective choice (i.e. in the drugged state, the subject experiences a given delay as longer than it remembered, causing a decrease in its preference for the Delayed lever). However, all drugs tested tended to shift preference towards the smaller reinforcer at high doses that significantly increased initiation latencies and omissions; thus this preference for the immediate reinforcer might be a non-specific drug effect. For example, a disinhibiting effect on operant behaviour, an impairment of stimulus control or an impairment of memory for the instrumental contingency resulting in delayed reward might all favour the response producing an immediate reinforcer, although it cannot be known which, if any, of these putative mechanisms were operating. Nevertheless, this general tendency makes the cue-dependent effect of amphetamine the more striking.

Effects of chlordiazepoxide

CDP was used as a positive control for possible non-specific drug effects on performance, because it does not affect the control over behaviour by conditioned reinforcers (Robbins *et al.*, 1983). As predicted by this account, it did not interact with the cue condition in determining choice of the two reinforcers. At the highest dose used, CDP reduced preference for the delayed reinforcer (increased impulsivity); this was true of high doses of all drugs used and may represent a non-specific drug effect (see above). At doses that did not severely disrupt responding (as assessed by the omission rate), an increase in impulsive responding was also observed, and at one dose CDP shifted preferences in both directions within the session (3.2 mg/kg, Cue group), being the only occasion when it caused a decrease in impulsivity. CDP had no consistent effects on latencies (other than at the highest dose) or on nose-poking during the delay.

The finding that CDP generally reduced tolerance of delayed reward is in contrast to the demonstration by Evenden and Ryan (1996) that another benzodiazepine, diazepam, increased preference for the delayed reward in this task. However, the present finding is in accord with the effect of CDP and other benzodiazepines to promote an 'impulsive' strategy in a T-maze task (Thiébot *et al.*, 1985). The action of benzodiazepines to increase impulsivity has been suggested to depend on a decrease in serotonin neurotransmission (Bizot & Thiébot, 1996; Bizot *et al.*, 1999); indeed, one benzodiazepine subgroup, the triazolobenzodiazepines, can increase 5-HT release and has been shown to reduce impulsive choice in the T-maze (Bizot *et al.*, 1999). While CDP blocks serotonin release *in vivo* (Soubrié *et al.*, 1983), and serotonin depletion has been shown to increase impulsivity in some tasks (Wogar *et al.*, 1993b; Richards & Seiden, 1995) with serotonin reuptake inhibitors having the opposite effect (Thiébot, 1986), this does not readily explain the discrepancy: diazepam and chlordiazepoxide have similar effects *in vitro* on midbrain

serotonin neurons (Thiébot *et al.*, 1982), and Evenden and Ryan (1996) found that the mixed serotonin receptor antagonist metergoline decreased impulsivity in the present task. While the effects of benzodiazepines on impulsive behaviour and the basis of these effects remain uncertain, the present results suggest that signals during a delay to reinforcement do not contribute to their action.

Effects of alpha-flupenthixol

In general, doses of α -flupenthixol that did not severely disrupt responding had small effects to reduce preference for the large, delayed reinforcer (i.e. to reduce tolerance of delay or promote impulsive choice). Its effects in the Cue condition were therefore opposite to those of amphetamine, as was predicted from its action as a dopamine receptor antagonist. Although interactions with the cue were not marked, those interactions were in the predicted direction: α -flupenthixol had a greater capacity to reduce tolerance of delay when the cue was present. As dopamine receptor antagonists, including α -flupenthixol, tend to impair the control over behaviour by conditioned reinforcers and its potentiation by amphetamine (Robbins *et al.*, 1983; Cador *et al.*, 1991; Wolterink *et al.*, 1993; Killcross *et al.*, 1997a), these results are consistent with the conditioned reinforcement hypothesis. Not only was α -flupenthixol able to impair the cue's effects to support choice of the large reinforcer, but it dose-dependently abolished the ability of the cue to sustain nose-poking in the food magazine during the delay (a form of conditioned approach behaviour). Taken together with the amphetamine result, this suggests that dopamine-dependent mechanisms contribute to the capability to choose a delayed reward by contributing to the effectiveness of conditioned reinforcers. However, α -flupenthixol also promoted impulsive choice in the absence of the cue; as this was an effect common to all three drugs tested, this may represent a non-specific disinhibiting effect or lack of stimulus control, such as has been observed for other neuroleptic drugs (Canon, 1979; Szostak & Tombaugh, 1981).

Conclusions

One function of conditioned reinforcement is to bridge temporal gaps between an animal's actions and primary reinforcement. This capacity can assist animals in learning discriminations based on delayed reinforcement (Grice, 1948), but can also contribute to performance of well-learned tasks. In artificial situations, conditioned reinforcers can even control behaviour to the detriment of performance (Williams & Dunn, 1991). The present study has demonstrated that stimuli present during a delay to reinforcement, probably by acting as conditioned reinforcers, can influence the effects of psychomotor stimulants. This has implications for the understanding and treatment of disorders of impulsive choice in humans, including ADHD; in particular, it suggests that the maximum benefit of psychostimulant treatment in this disorder will be obtained when behaviour is highly controlled by conditioned reinforcers, and when the availability of delayed reward is clearly signalled (see also Sagvolden *et al.*, 1998). In addition, it supports the idea that 'delay discounting' of the efficacy of future rewards is not a unitary process (Ainslie, 1975), but rather that the observed phenomenon of discounting arises from several underlying processes, of which conditioned reinforcement is one.

Chapter 7.

Contributions of limbic and prefrontal circuitry to choice of delayed reinforcement

Abstract. Impulsive choice, the inability to choose a large delayed reward in preference to an immediate but small reward, is an important but poorly-understood phenomenon. As impulsive choice may result from an insensitivity to delayed reinforcement, and limbic corticostriatal circuits have been implicated in reinforcement processes, the present experiments investigated the contribution of components of the prefrontal cortex and ventral striatum to rats' ability to choose a delayed reward. Rats were trained on a two-lever discrete-trial delayed reinforcement task in which they chose one food pellet delivered immediately or four pellets delivered after a delay; this delay increased from 0 to 60 s during each session. Subjects developed a characteristic within-session shift in preference, choosing the larger reinforcer at short delays, but the smaller reinforcer when the delay was long. Once trained, the rats were assigned to matched groups and received excitotoxic lesions of the perigenual anterior cingulate cortex (ACC), medial prefrontal cortex (mPFC), or nucleus accumbens core (AcbC); they were then retested. Lesions of the ACC had no effect on subjects' capacity to choose the delayed reward, or their ability to respond to removal of the delays by choosing the large reward consistently, though ACC-lesioned subjects were slower to collect the larger reward than sham-operated controls. Lesions of the mPFC induced a 'flattening' of the within-session shift in preference, but subjects still responded normally to removal of the delays, suggesting a loss of temporal stimulus control. Lesions of the AcbC dramatically and persistently impaired subjects' ability to choose the large reinforcer when it was delayed, even though subjects discriminated the two reinforcers. It is suggested that dysfunction of the AcbC may be a key element in the pathology of impulsivity. In a different version of the task, intra-accumbens amphetamine was found to have slight but inconsistent effects to reduce preference for the delayed reinforcer, though this effect did not depend on whether the delayed reward was signalled or unsignalled.

INTRODUCTION

Impulsive choice is exemplified by the tendency of an individual to choose a reward that is small, poor, or ultimately disastrous, but is available immediately, in preference to a larger reward that is only obtainable after a period of time (Ainslie, 1975). Impulsive choice may reflect reduced efficacy of delayed reinforcement. It has been considered a normal human characteristic (Aristotle, 350 BC / 1925), but impulsive choice contributes to deleterious states such as drug addiction (Poulos *et al.*, 1995; Heyman, 1996; Bickel *et al.*, 1999; Evenden, 1999a; Mitchell, 1999) and has been suggested to underlie a number of other clinical disorders, including attention-deficit/hyperactivity disorder (ADHD; Sagvolden *et al.*, 1998; Sagvolden & Sergeant, 1998).

Little is known of the neuroanatomical basis of impulsive choice. However, three lines of evidence suggest the nucleus accumbens (Acb) and its cortical afferents, including the anterior cingulate and me-

dial prefrontal cortices (ACC, mPFC), as candidate structures that may be involved in regulating choice between alternative reinforcers.

First, these structures have been firmly implicated in reinforcement processes. The Acb, once suggested to mediate the reinforcing efficacy of natural and artificial rewards (see Koob, 1992) (and also Wise, 1981; 1982; 1985; 1994), is now thought not to be necessary for this, but instead to be a key site for the motivational impact of impending rewards (reviewed by Robbins & Everitt, 1996; Salamone *et al.*, 1997; Everitt *et al.*, 1999; Parkinson *et al.*, 2000a). Many of its afferents have also been shown to be involved in reward-related learning, including the ACC (Chapter 3; Bussey *et al.*, 1997a; Bussey *et al.*, 1997b; Parkinson *et al.*, 2000c) and mPFC (e.g. Balleine & Dickinson, 1998a; Richardson & Gratton, 1998; Bechara *et al.*, 1999; Tzschentke, 2000).

Second, these regions are important recipients of dopaminergic and serotonergic afferents (Fallon & Loughlin, 1995; Halliday *et al.*, 1995), and pharmacological manipulations of dopamine and serotonin systems have been shown to affect impulsive choice in rats (Sagvolden *et al.*, 1992; Wogar *et al.*, 1993b; Richards & Seiden, 1995; Charrier & Thiébot, 1996; Evenden & Ryan, 1996; Richards *et al.*, 1997a; Evenden, 1998; Bizot *et al.*, 1999; Evenden, 1999b; Evenden & Ryan, 1999; Ho *et al.*, 1999; Richards *et al.*, 1999; Cardinal *et al.*, 2000b; Wade *et al.*, 2000).

Third, abnormalities of these regions have been detected in humans with ADHD, and in animal models of ADHD. Abnormal functioning of prefrontal cortical regions, including medial prefrontal and anterior cingulate cortex, has been observed in ADHD patients (Ernst *et al.*, 1998; Bush *et al.*, 1999; Rubia *et al.*, 1999). In the spontaneously hypertensive rat (SHR), widely used as an animal model of ADHD (Wultz *et al.*, 1990; Sagvolden *et al.*, 1992; Sagvolden *et al.*, 1993; Sagvolden, 2000), differences in dopamine receptor density and gene expression have been observed within the core and shell regions of the Acb (Papa *et al.*, 1996; Carey *et al.*, 1998; Papa *et al.*, 1998; Sadile, 2000). Abnormalities of dopamine release have been detected in the Acb (de Villiers *et al.*, 1995; Russell *et al.*, 1998; Russell, 2000) and prefrontal cortex (Russell *et al.*, 1995), in addition to possible dysfunction in the dorsal striatum and amygdala (Russell *et al.*, 1995; Papa *et al.*, 2000).

Evenden and Ryan (1996) developed a model of impulsive choice in which food-restricted rats choose between a small, immediate reward and a large, delayed reward in discrete trials, the delay to the large reinforcer being increased in steps as the session progressed. The present study investigated the effects of excitotoxic lesions of the ACC, mPFC, and AcbC on performance of a modified version of this task. Potentially, the lesions might affect learning of the task; in order to avoid this confounding factor, subjects were trained before the lesions were made. As it was demonstrated in Chapter 6 that explicit signals present during a delay to reinforcement may affect the response to a behavioural or pharmacological manipulation, the simplest situation was used, with no signals present during the delay to reinforcement. After subjects had been tested post-operatively, all delays were removed from the task to establish whether lesioned subjects remained sensitive to the delays.

Finally, an experiment was conducted to investigate the role of the Acb in the effects of amphetamine on impulsive choice. Amphetamine was injected directly into the Acb before animals chose between a small, immediate reward and a large, delayed reward in discrete trials. As the effects of amphetamine depend in part upon signals present during the delay to reward (Chapter 6; Cardinal *et al.*, 2000b), intra-Acb amphetamine was administered to two groups of subjects, trained with or without a cue stimulus present during this delay. As discussed in Chapter 6 (p. 192), it was anticipated that intra-Acb amphetamine would enhance the conditioned reinforcing properties of such a stimulus, promoting choice of the delayed reward in the cued group.

EXPERIMENT 1. EFFECTS OF LESIONS OF THE ANTERIOR CINGULATE CORTEX

Methods

Twenty-four naïve rats were maintained at 90% of their free-feeding mass and trained on the same delay-of-reinforcement task used in Chapter 6 (*q.v.*). They were first trained to press levers for sucrose pellets. (In Chapter 6, subjects were allowed to respond freely on an FR1 schedule on the left lever until they had acquired at least 50 reinforcers in 30 minutes, and then trained on the right lever, with no limit on the number of reinforcers available in each 30-min session. However, it was observed that subjects tended to acquire responding more rapidly, and thus accrue more reinforcers, on the lever trained second; thus, for all studies in the present chapter, subjects were trained until they had accrued an *overall total* of 50 reinforcers on each lever in turn; when this limit had been reached, the lever was retracted and the session finished.) Next, they were trained to nosepoke to initiate discrete-trial presentations of the levers, before being trained on the main delay-of-reinforcement task for 19 sessions. No cues were present during the delays to reinforcement. After this, they were assigned to matched groups by ranking all subjects according to the regression slope measure (see Chapter 6, p. 174), calculated using data from the last 3 pre-operative sessions. The ranked list was divided into pairs, and from each pair one subject was assigned to the sham group and the other to the ACCX group, at random. It was subsequently ensured that both groups had achieved criterion sensitivity to delay (see Chapter 6), and that there were no significant pre-operative differences in the absolute level of preference.

Subjects then received lesions of the anterior cingulate cortex (ACCX, $n = 12$) or sham lesions (sham, $n = 12$). At the time of surgery, their body mass was 329–379 g. Following recovery, they were retested on the basic task for 7 sessions to obtain a baseline of performance. After this, 4 sessions were given in which all delays were omitted in alternate sessions (DNDN design; D = delays present, N = no delays). Half of the subjects began this test with the delays present, and half with no delays (counterbalanced across groups).

Results

Histology

One subject in the ACCX group died post-operatively (subject H11). Histological analysis revealed that the lesion was incomplete in two subjects (subjects H2, H22), who were excluded, leaving 9 in the ACCX group (H1, H3, H8, H10, H14, H15, H18, H19, H21) and 12 in the sham group (H4, H5, H6, H7, H9, H12, H13, H16, H17, H20, H23, H24). Neuronal loss and associated gliosis in the lesion group extended from ~3.0 mm anterior to bregma to ~0.3 mm posterior to bregma, damaging perigenual Cg1 and Cg2, and in some cases Cg1 more anteriorly. There was no damage to PrL, IL, PCC, or the corpus callosum. Within the ACCX group, there was some heterogeneity; 5 of these animals had lesions encompassing the entire ventral perigenual region, including the ventral portion of Cg2 at 1.6–1.7 mm anterior to bregma (H1, H10, H14, H15, H19; see Figure 82), while 4 did not (H3, H8, H18, H21; see Figure 83). Representative photomicrographs of ACC lesions were shown in Chapter 3 (p. 80).

Anterior cingulate cortex: schematic of lesions

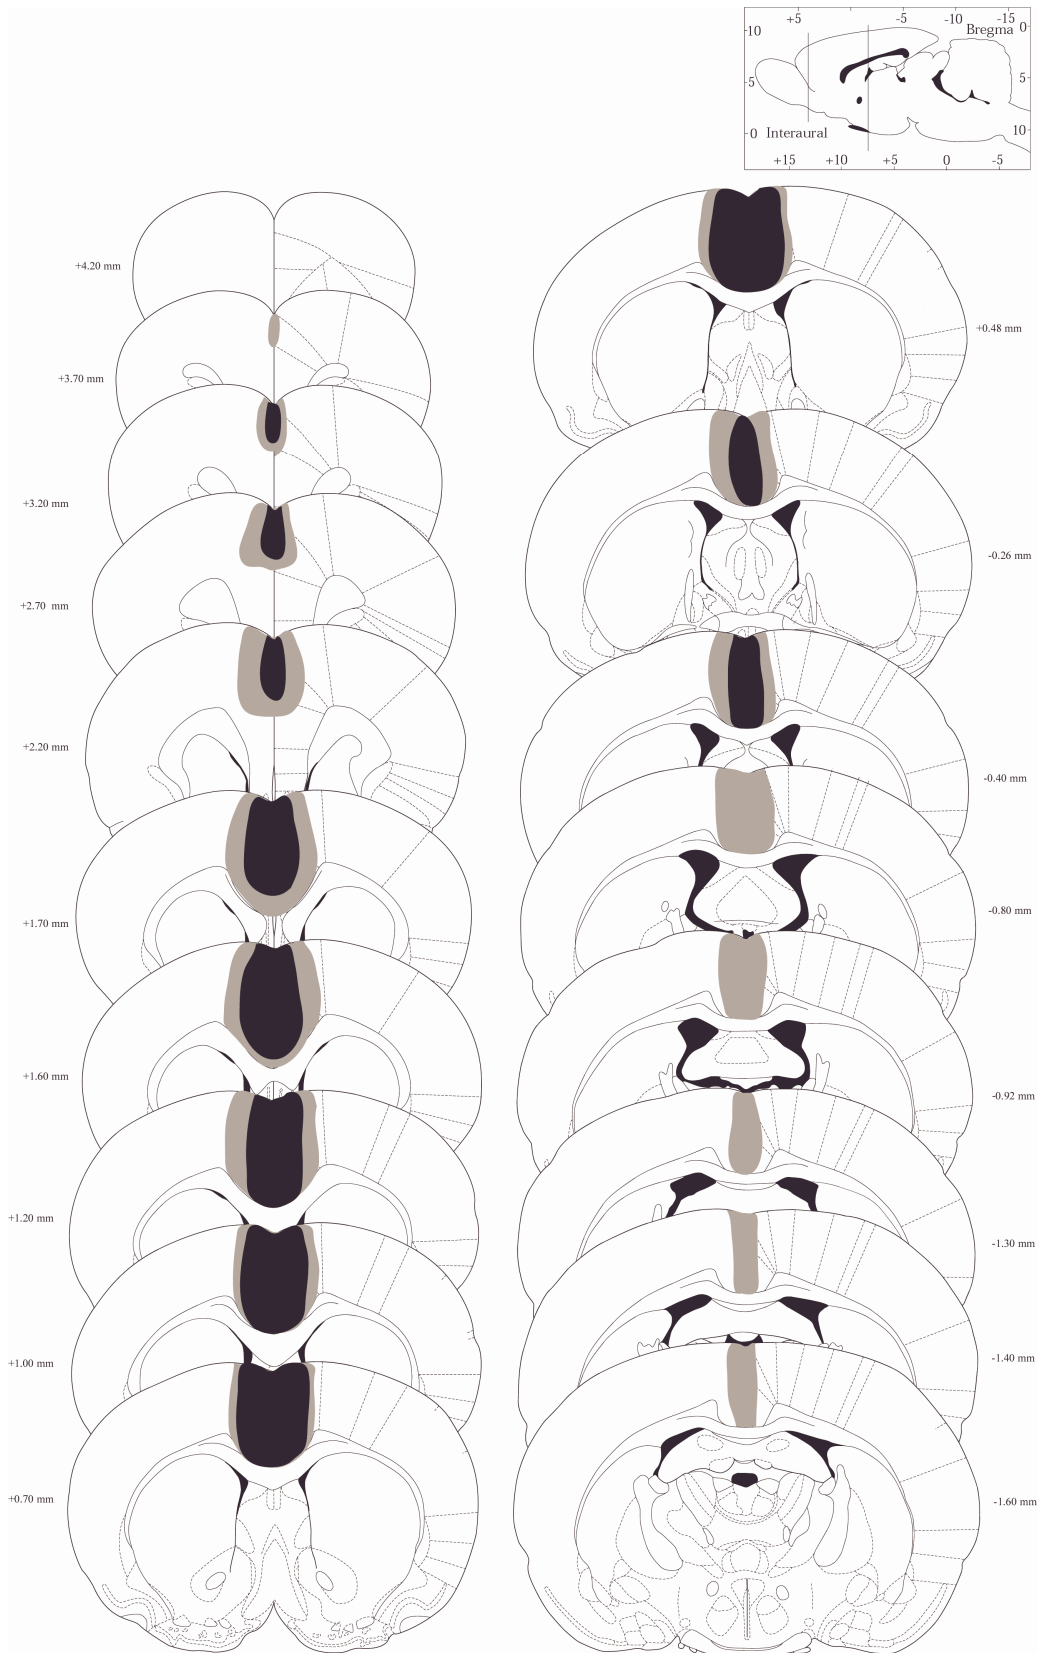


Figure 82. Lesions of the ACC, including the ventral perigenual region. Subjects were classified as having whole or partial ACC lesions on the basis of whether the ventral portion of Cg2 in the 'cup' of the genu was lesioned (seen here in sections +1.6 and +1.7 mm from bregma). Grey shading indicates the extent of the largest area of neuronal loss, and black the smallest. Diagrams are taken from Paxinos & Watson (1998). (Subjects: H1, H10, H14, H15, H19.)

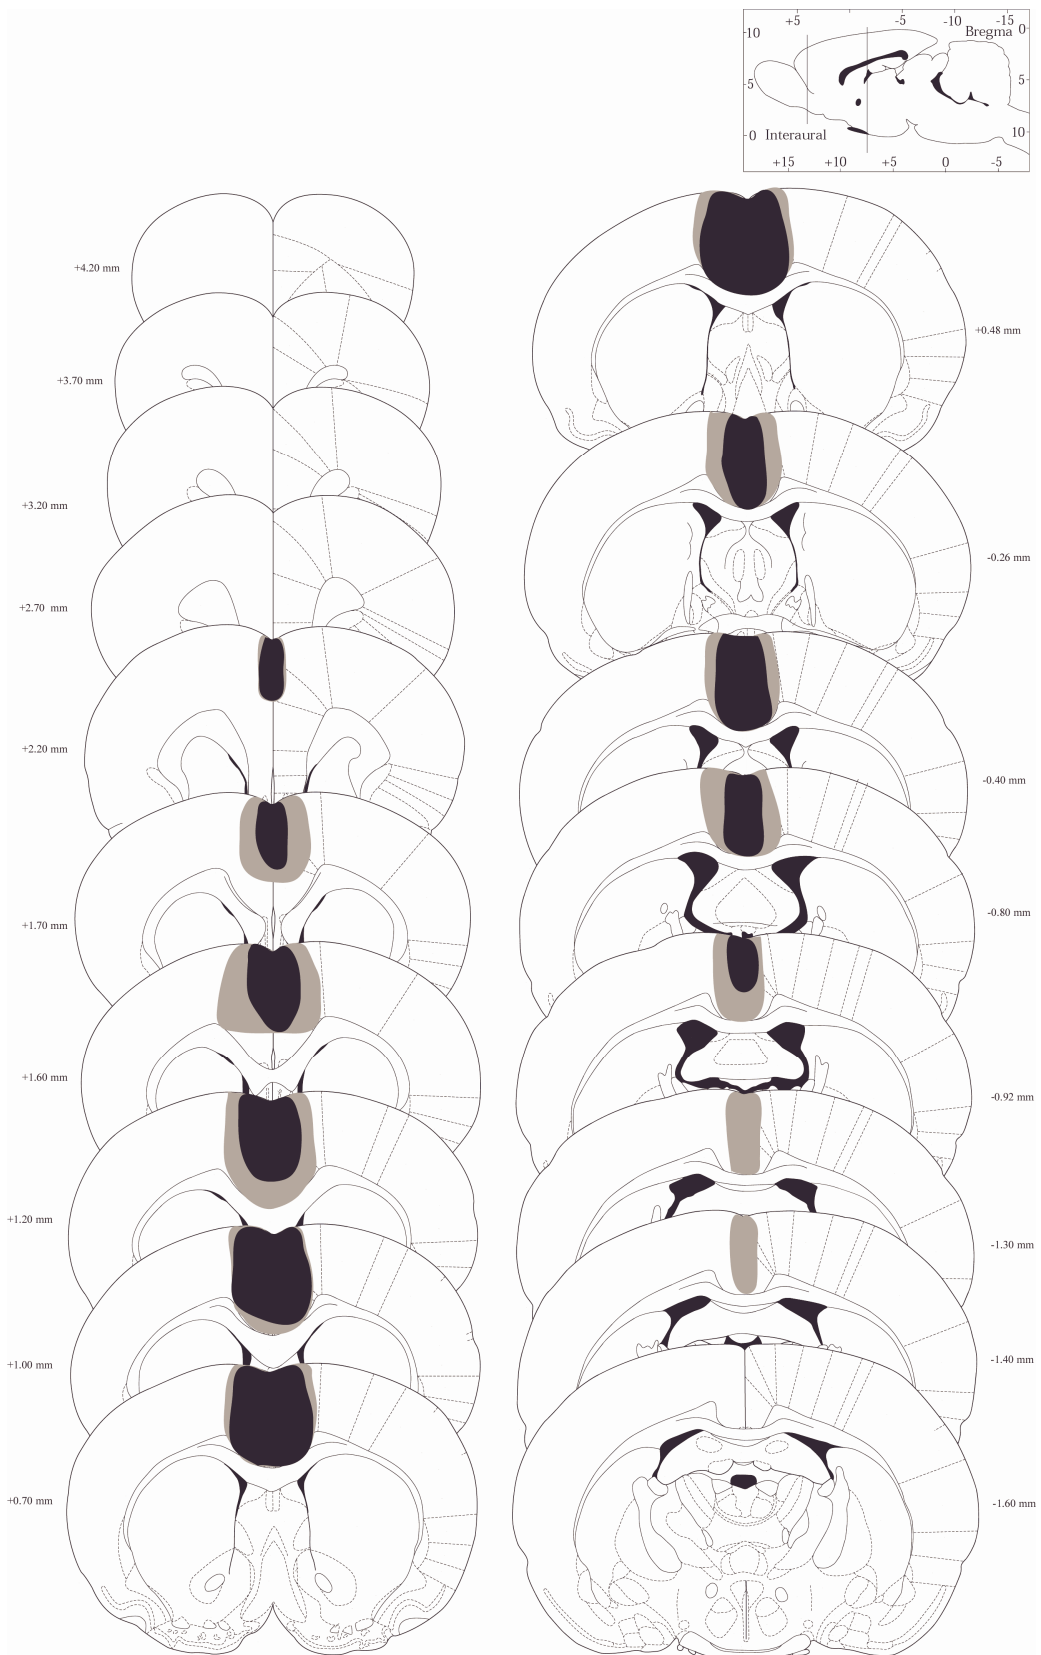


Figure 83. Lesions of the ACC, *excluding* the ventral perigenual region (compare Figure 82). (Subjects: H3, H8, H18, H21.)

Pre-operative acquisition

The groups remained matched after histological selection: there were no differences in the pre-operative pattern of choice (Figure 84A). Choice ratios from the last 3 pre-operative sessions were analysed using the model $\text{group} \times (\text{delay} \times \text{S})$. While there was a highly significant effect of delay ($F_{1,858,35,299} = 44.349$, $\tilde{\epsilon} = .464$, $p < .001$), there was no effect of group and no group \times delay interaction ($F_s < 1$, NS).

Baseline post-operative performance

Choice. There were no differences between sham and ACCX groups in the pattern of choice observed for the 7 baseline sessions (Figure 84B–D). Analysis of choice ratios demonstrated that the effect of delay remained highly significant ($F_{2,404,45,684} = 53.46$, $\tilde{\epsilon} = .601$, $p < .001$), but there was no effect of group and no group \times delay interaction ($F_s < 1$, NS). The rapidity of the within-session shift in preference, as assessed by the slope measure, did not differ either, and did not alter across the post-operative sessions; analysis using the model $\text{group}_2 \times (\text{session}_7 \times \text{S})$ revealed no effect of any term ($F_s < 1.101$, NS). A separate comparison between shams and the subgroup of ACC-lesioned animals with complete ventral perigenual damage did not alter these conclusions (all terms involving group, $F_s < 1$, NS).

Omissions. Responding was reliable, with all animals regularly sampling both levers, and the two groups did not differ in the number of omissions made. An analysis of the percentage of trials on which an omission occurred, across all delays, revealed no effect of group ($F_{1,19} = 2.686$, NS).

Initiation latency. While initiation latencies increased with delay (from 1.20 ± 0.06 s at zero delay to 1.55 ± 0.12 s at the maximum delay), there were no differences between the two groups (delay: $F_{2,639} = 11.572$, $\tilde{\epsilon} = .66$, $p < .001$; group and delay \times group, $F_s < 1$, NS).

Choice latency. Subjects responded faster on the lever that produced the larger reward, particularly at short delays (mean latencies at zero delay: large reward 0.96 ± 0.05 s, small reward 1.34 ± 0.014 s; at 60 s delay: large reward 0.97 ± 0.05 s, small reward 1.01 ± 0.04 s). However, there were no group differences. An analysis using the model $\text{group} \times (\text{response} \times \text{delay} \times \text{S})$ revealed a response \times delay interaction ($F_{2,577,43,804} = 5.676$, $\tilde{\epsilon} = .644$, $p = .003$), as well as main effects of response ($F_{1,17} = 5.195$, $p = .036$) and delay ($F_{2,021,34,357} = 4.997$, $\tilde{\epsilon} = .505$, $p = .012$). However, no terms involving group were significant ($F_s < 1$, NS).

Collection latency. Lesioned subjects were slower to collect the larger reward (Figure 84D). An analysis by group \times (response \times delay \times S) revealed a group \times response interaction ($F_{1,17} = 15.1$, $p = .001$) as well as a main effect of delay ($F_{2,794,47,501} = 3.042$, $\tilde{\epsilon} = .699$, $p = .041$), reflecting slightly longer collection latencies at long delays. No other terms were significant (closest to significance: response, $F_{1,17} = 2.466$, $p = .135$).

Nosepoking during the delay. While there was a small tendency for subjects to spend a greater proportion of time nose-poking at longer delays, no group differences were found. An analysis of the percentage of the delay spent nose-poking, using the model $\text{group}_2 \times (\text{delay}_4 \times \text{S})$, revealed an effect of delay ($F_{2,13,40,463} = 3.36$, $\tilde{\epsilon} = .71$, $p = .042$) but no effect of group ($F_{1,19} = 1.663$, NS) and no group \times delay interaction ($F_{2,13,40,463} = 1.938$, $\tilde{\epsilon} = .71$, NS).

Effect of omitting all delays

Both groups remained sensitive to the removal of delays, shifting their preference towards the large reinforcer under these conditions (Figure 84E). Analysis of choice ratios using the model $\text{group}_2 \times (\{\text{Delays versus No Delays}\}_2 \times \text{trial block}_5 \times \text{S})$ revealed a highly significant interaction between the Delay/No Delay factor and the trial block ($F_{3,302,62,731} = 39.346$, $\tilde{\epsilon} = .825$, $p < .001$), in addition to main effects of the

Delay/No Delay factor ($F_{1,19} = 30.235, p < .001$) and the trial block ($F_{4,76} = 33.679, p < .001$), However, there were no group differences (terms involving group: $F_s < 1.392$, NS). As in the previous study (Chapter 6), a significant shift of preference persisted in the absence of delays (simple effect of trial block in the No Delay condition: $F_{4,76} = 2.542, p = .046$), although it was slight.

Summary

Lesions of the ACC did not affect subjects' ability to choose a delayed reward; their pattern of choice was indistinguishable from that of sham-operated controls, and their behaviour remained sensitive to removal of the delays. The only behavioural difference apparent was that ACC-lesioned subjects collected the large reward somewhat slower than controls.




Figure 84. Effects of lesions of the ACC on performance of the delayed-reinforcement choice task. **A:** Pre-operative performance — data from the last 3 sessions preceding surgery. **B:** Post-operative performance — data from the first 7 sessions following surgery. **C:** Slope measures before and after surgery. This slope measure is the linear regression of %choice of the large reinforcer against $\log(\text{delay} + 1 \text{ s})$, calculated for each session. More negative slopes indicate a larger within-session shift from the large to the small reinforcer as the delay lengthens. **D:** Latencies to collect reward post-operatively, averaged across all delays. ACC-lesioned rats were slower to collect the large reward. **E:** Effect of omitting all delays in alternating sessions (2SED, twice the standard error of the difference for the three-way interaction).

EXPERIMENT 2. EFFECTS OF LESIONS OF MEDIAL PREFRONTAL CORTEX

Methods

Twenty-four naïve subjects were trained and assigned to two groups as in Experiment 1 (p. 197). They then received lesions of the medial prefrontal cortex (mPFC group, $n = 14$) or sham lesions ($n = 10$). At the time of surgery, they weighed 276–373 g. Following recovery, they were retested on the basic task for 7 sessions. After this, 4 sessions were given in which all delays were omitted in alternate sessions (ABAB design), as before. Finally, a 2-h locomotor test was given.


Results

Histology


There were no postoperative deaths. One rat was excluded from the mPFC group because its lesion was unilateral (M3), and two because the lesion extended beyond the genu posteriorly (M4, M7), leaving 11 in the mPFC group (M6, M9, M10, M12, M13, M14, M16, M18, M20, M22, M24) and 10 in the sham group (M1, M2, M5, M8, M11, M15, M17, M19, M21, M23). Within the mPFC group, neuronal loss and associated gliosis extended from approximately 5.0 to 1.7 mm anterior to bregma. Within this region, there was extensive damage to prelimbic cortex, with damage also occurring in infralimbic cortex, dorsal Cg1, and medial orbital cortex. There was no damage posterior to the genu. Representative photomicrographs are shown in Figure 85, and schematics (indicating the largest and smallest extent of the lesions) are shown in Figure 86.

Medial prefrontal cortex: photomicrographs


A. Sham-operated rat (low power)



B. Sham-operated rat (high power)



C. mPFC-lesioned rat (low power)



D. mPFC-lesioned rat (high power)




Figure 85. Lesions of the mPFC: photomicrographs of sections at approximately 2.6 mm anterior to bregma, stained with cresyl violet. **A & B:** sham-operated rat (M2, secondary motor cortex; Cg1, cingulate area 1; PrL, prelimbic cortex; IL, infralimbic cortex; fmi, forceps minor of the corpus callosum). **C & D:** mPFC-lesioned rat. Dotted lines show the extent of the lesion. **Left-hand panels** are low-magnification views (scale bars are 1 mm); **right-hand panels** are high-magnification views (scale bars are 0.1 mm). Arrowheads indicate the position of identical structures in corresponding pairs of photomicrographs.

Medial prefrontal cortex: schematic of lesions

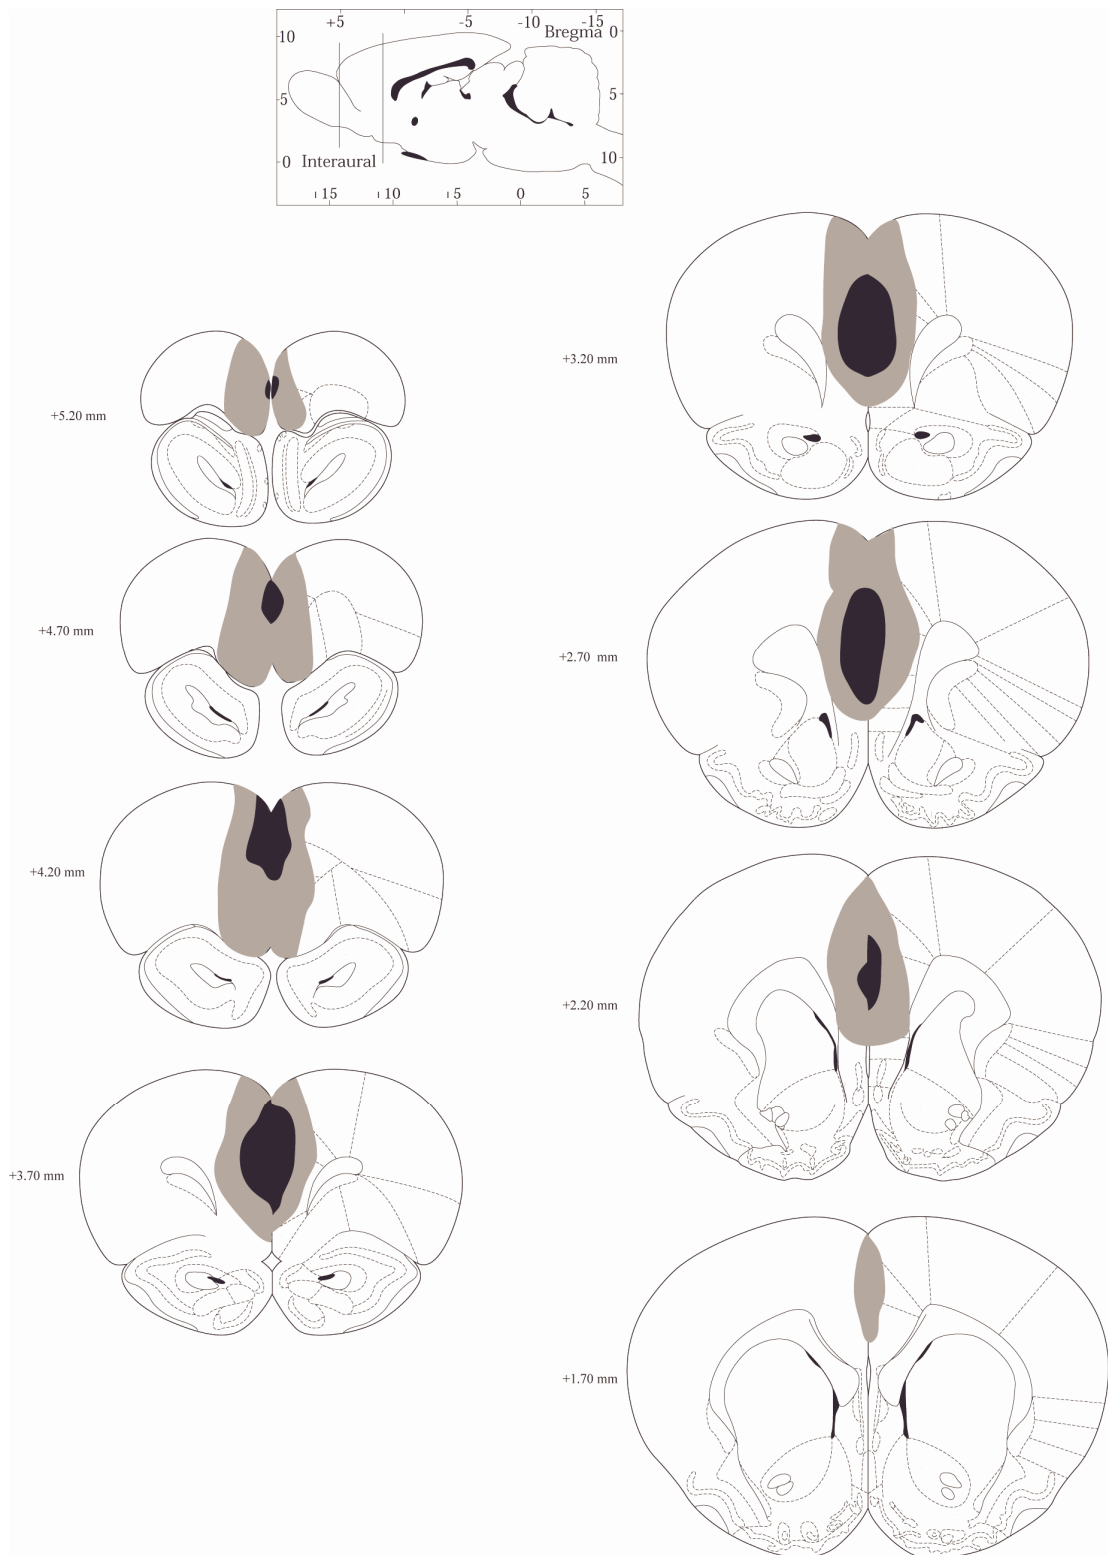


Figure 86. Lesions of the mPFC (subjects M6, M9, M10, M12, M13, M14, M16, M18, M20, M22, M24). Grey shading indicates the extent of the largest area of neuronal loss, and black the smallest. Diagrams are taken from Paxinos & Watson (1998).

Body mass

The two groups did not differ in body mass, either at the start or the end of behavioural testing ($F_s < 1$, NS).

Pre-operative acquisition

The groups remained matched after histological selection: there were no differences in the pre-operative pattern of choice (Figure 88A, p. 208). Choice ratios calculated from the last 3 pre-operative sessions were analysed using the model group \times (delay \times S). While there was a highly significant effect of delay ($F_{2.528,48.033} = 18.429$, $\tilde{\epsilon} = .632$, $p < .001$), there was no effect of group and no group \times delay interaction ($F_s < 1$, NS).

Baseline post-operative performance

Choice. Although the mPFC-lesioned group exhibited a within-session shift in preference, this shift was less pronounced than in the sham group (Figure 88B, p. 208). Analysis of choice ratios using the model group \times (delay \times S) revealed a group \times delay interaction ($F_{2.891,54.926} = 3.188$, $\tilde{\epsilon} = .723$, $p = .032$), in addition to a main effect of delay ($F_{2.891,54.926} = 26.831$, $\tilde{\epsilon} = .723$, $p < .001$). There was no main effect of group ($F < 1$, NS). Separate analyses of each group demonstrated that both groups shifted their preference from the large to the small reinforcer as the delay increased (effect of delay in the sham group: $F_{2.819,25.367} = 17.499$, $\tilde{\epsilon} = .705$, $p < .001$; in the mPFC group: $F_{4,40} = 8.87$, $p < .001$). At no *individual* delay was preference different between the two groups (simple effects of group at each delay: $F_s < 1.218$, NS).

This interpretation was confirmed by analysis of the regression slope measure, which was substantially higher post-operatively in the mPFC group, indicating a flattened within-session shift in preference (Figure 88C; more negative values of this measure indicate a more pronounced shift from large to small reinforcer across the session). Analysis of slope measures from pre- and post-operative sessions, using the model group \times (session \times S), revealed a highly significant group \times session interaction ($F_{16.491,313.333} = 2.286$, $\tilde{\epsilon} = .66$, $p = .003$), in addition to a main effect of session ($F_{16.491,313.333} = 6.265$, $\tilde{\epsilon} = .66$, $p < .001$); there was no main effect of group in this analysis ($F < 1$, NS). This interaction was not due to pre-operative differences between the groups: analysis of pre-operative sessions 1–19 revealed a main effect of session ($F_{9,914,188.366} = 7.375$, $\tilde{\epsilon} = .551$, $p < .001$), but no effect of group ($F_{1,19} = 1.288$, NS) and no group \times session interaction ($F_{9,914,188.366} = 1.399$, $\tilde{\epsilon} = .551$, NS). Post-operatively, however, slope measures were substantially higher (less negative) in the mPFC group, with analysis of post-operative sessions 20–26 revealing a main effect of group ($F_{1,19} = 4.848$, $p = .04$). This pattern did not change during post-operative testing: there was no effect of session ($F_{6,114} = 1.127$, NS) and no group \times session interaction ($F < 1$, NS).

Omissions. Responding was reliable, with all animals regularly sampling both levers, and the two groups did not differ in the number of omissions made. An analysis of the percentage of trials on which an omission occurred, across all delays, revealed no effect of group ($F < 1$, NS).

Initiation latency. Subjects were slower to initiate trials as the session progressed and the delays lengthened, but there were no differences between mPFC and sham groups in this respect. Analysis of initiation latencies using the model group \times (delay \times S) revealed a main effect of delay ($F_{2.456,46.657} = 4.637$, $\tilde{\epsilon} = 4.637$, $p = .01$) but no significant terms involving group (group: $F_{1,19} = 2.255$, NS; group \times delay: $F < 1$, NS).

Choice latency. Lesioned rats were slower to respond on the levers (Figure 88D), and initiation latencies were generally longer for all subjects at the start of the session. Analysis using the model group \times

(delay \times response \times S) demonstrated main effects of group ($F_{1,16} = 7.741, p = .013$) and delay ($F_{2,246,35.934} = 5.137, \tilde{\epsilon} = .561, p = .009$), but no other significant terms (response \times delay: $F_{2,373,37.975} = 2.119, \tilde{\epsilon} = .593, NS$; other terms: $F < 1, NS$).

Collection latency. The lesion did not affect collection latencies. Subjects collected the immediate reward slightly faster than the delayed reward, and were slower to collect rewards as the session progressed; these two tendencies were statistically independent. Analysis using the model group \times (delay \times response \times S) showed main effects of response ($F_{1,16} = 6.305, p = .023$) and delay ($F_{3,663,58.601} = 2.647, \tilde{\epsilon} = .916, p = .047$), but no other significant terms ($F_s < 1.823, NS$).

Nosepoking during the delay. The lesion did not affect nosepoking behaviour, and nosepoking occurred at a constant rate at all delays. Analysis using the model group \times (delay \times S) revealed no effect of any term (group: $F_{1,17} = 2.14, NS$; other terms: $F < 1.454, NS$).

Effect of omitting all delays

Both groups remained sensitive to the removal of delays, shifting their preference towards the large reinforcer under these conditions (Figure 88E). Analysis of choice ratios using the model group₂ \times ({Delays versus No Delays}₂ \times trial block₅ \times S) revealed a highly significant interaction between the Delay/No Delay factor and the trial block ($F_{2,269,43.116} = 29.442, \tilde{\epsilon} = .567, p < .001$) in addition to main effects of the Delay/No Delay factor ($F_{1,19} = 22.949, p < .001$) and of trial block ($F_{3,253,61.813} = 17.117, \tilde{\epsilon} = .813, p < .001$), but there were no significant terms involving group ($F_s < 1.693, NS$).

Locomotor activity in a novel environment

The mPFC group were not significantly hyperactive (Figure 87). Following square-root transformation, analysis of the total number of infrared beam interruptions using the model group₂ \times (bin₁₂ \times S) revealed an effect of bin ($F_{6,048,114.909} = 16.046, \tilde{\epsilon} = .55, p < .001$), reflecting habituation, but no other significant term (group: $F_{1,19} = 2.168, NS$; group \times bin: $F < 1, NS$).




Figure 87. Locomotor activity in a novel environment (120-min session scored in 10-min bins). There were no significant differences between the groups.

Summary

Lesions of the mPFC induced a 'flattening' of the normal within-session shift in preference from the large to the small reward, though lesioned subjects still exhibited this shift and remained sensitive to removal of the delays. They were also generally slower to respond on the levers.

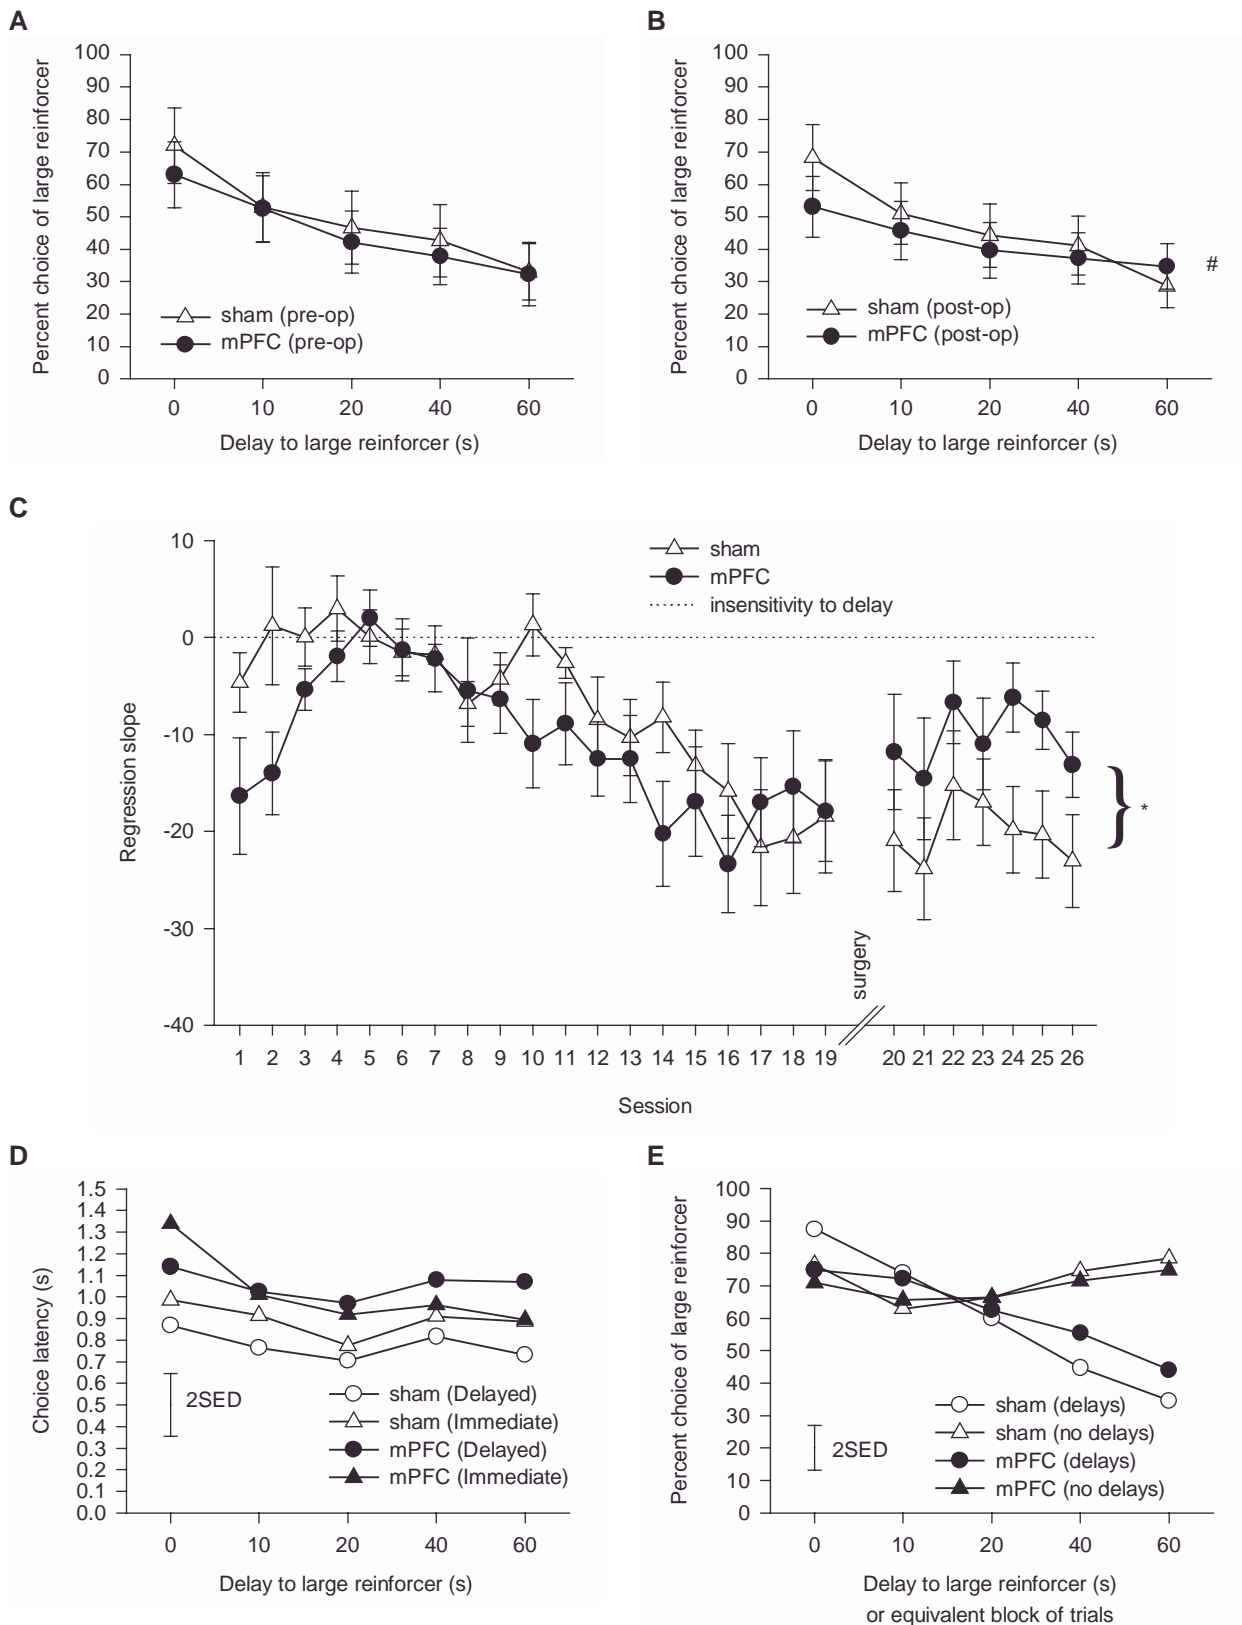


Figure 88. Performance of rats with lesions of the mPFC on the delayed-reinforcement choice task. **A:** Pre-operative performance — data from the last 3 sessions preceding surgery. **B:** Post-operative performance — data from the first 7 sessions following surgery (# $p < .05$, group \times delay interaction). **C:** Slope measures before and after surgery (* $p < .05$, post-operative difference between groups). **D:** Latencies to choose a lever; the mPFC group were significantly slower to respond. **E:** Effect of omitting all delays in alternating sessions. (2SED, twice the standard error of the difference for the relevant three-way interaction.)

EXPERIMENT 3. EFFECTS OF LESIONS OF THE NUCLEUS ACCUMBENS CORE

Methods

Twenty-four naïve subjects were trained and assigned to two groups as in Experiment 1 (p. 197). They then received lesions of the AcbC ($n = 14$) or sham lesions ($n = 10$). At the time of surgery, they weighed 315–372 g. Following recovery, they were retested on the basic task for 7 sessions, and given 4 sessions in which all delays were omitted in alternate sessions (ABAB design), as in Experiment 1.

As a deficit was observed during testing (before histological data were available), further behavioural tests were given to elucidate the nature of the deficit. First, the delay-omission test was repeated over 6 sessions, using an AAABBB design (three sessions with delays present, followed by three sessions with no delays, or vice versa). This test gave subjects longer to respond to the new contingencies. As before, half of the subjects began this test with the delays present, and half with no delays (counterbalanced across groups). Secondly, all animals were given a further 6 sessions with no delays, in an attempt to re-equalize the two groups' performance and ensure that all animals would come to prefer the lever producing the large reinforcer. Finally, the delays were re-introduced for a further 6 sessions.

Following completion of delayed reinforcement testing, subjects were given a 2-h locomotor test (methodologically identical to that used in Chapter 3, p. 78). After this, a pellet/chow consumption test was administered, as described below, before the animals were killed and perfused.

Food consumption tests

Food consumption was assessed using four tests, conducted in subjects' home cages (always with only one rat present) on separate days under conditions of food deprivation.

- (1) Subjects were given free access to 45-mg sucrose pellets (Rodent Diet Formula P, Noyes, Lancaster, NH) for 30 minutes; the amount eaten was recorded.
- (2) This test was repeated with the chow used as the maintenance diet.
- (3) The time taken to consume 50 sucrose pellets was recorded.
- (4) The time taken to consume an equivalent mass of chow (2.25 g) was recorded.

Results

Histology

There were no postoperative deaths. Histological analysis revealed that one subject in the core group (J23) had no damage to the Acb, one subject (J5) had an extensive lesion involving the septum, and two other subjects (J6, J19) had lesions encompassing a significant proportion of the AcbSh. These animals were excluded, leaving 10 subjects in the core group (J1, J11, J12, J13, J15, J16, J17, J18, J21, J24) and 10 in the sham group (J2, J3, J4, J7, J8, J9, J10, J14, J20, J22). Lesions of the AcbC encompassed most of the core subregion; neuronal loss and associated gliosis extended in an anteroposterior direction from approximately 2.5 mm to 0.5 mm anterior to bregma, and did not extend ventrally or caudally into the ventral pallidum or olfactory tubercle. Damage to the ventromedial caudate–putamen was occasionally seen; damage to AcbSh was restricted to the lateral edge of the dorsal shell. Schematics of the lesions are shown in Figure 89; representative photomicrographs of AcbC lesions were shown in Chapter 4 (p. 132).

Nucleus accumbens core: schematic of lesions

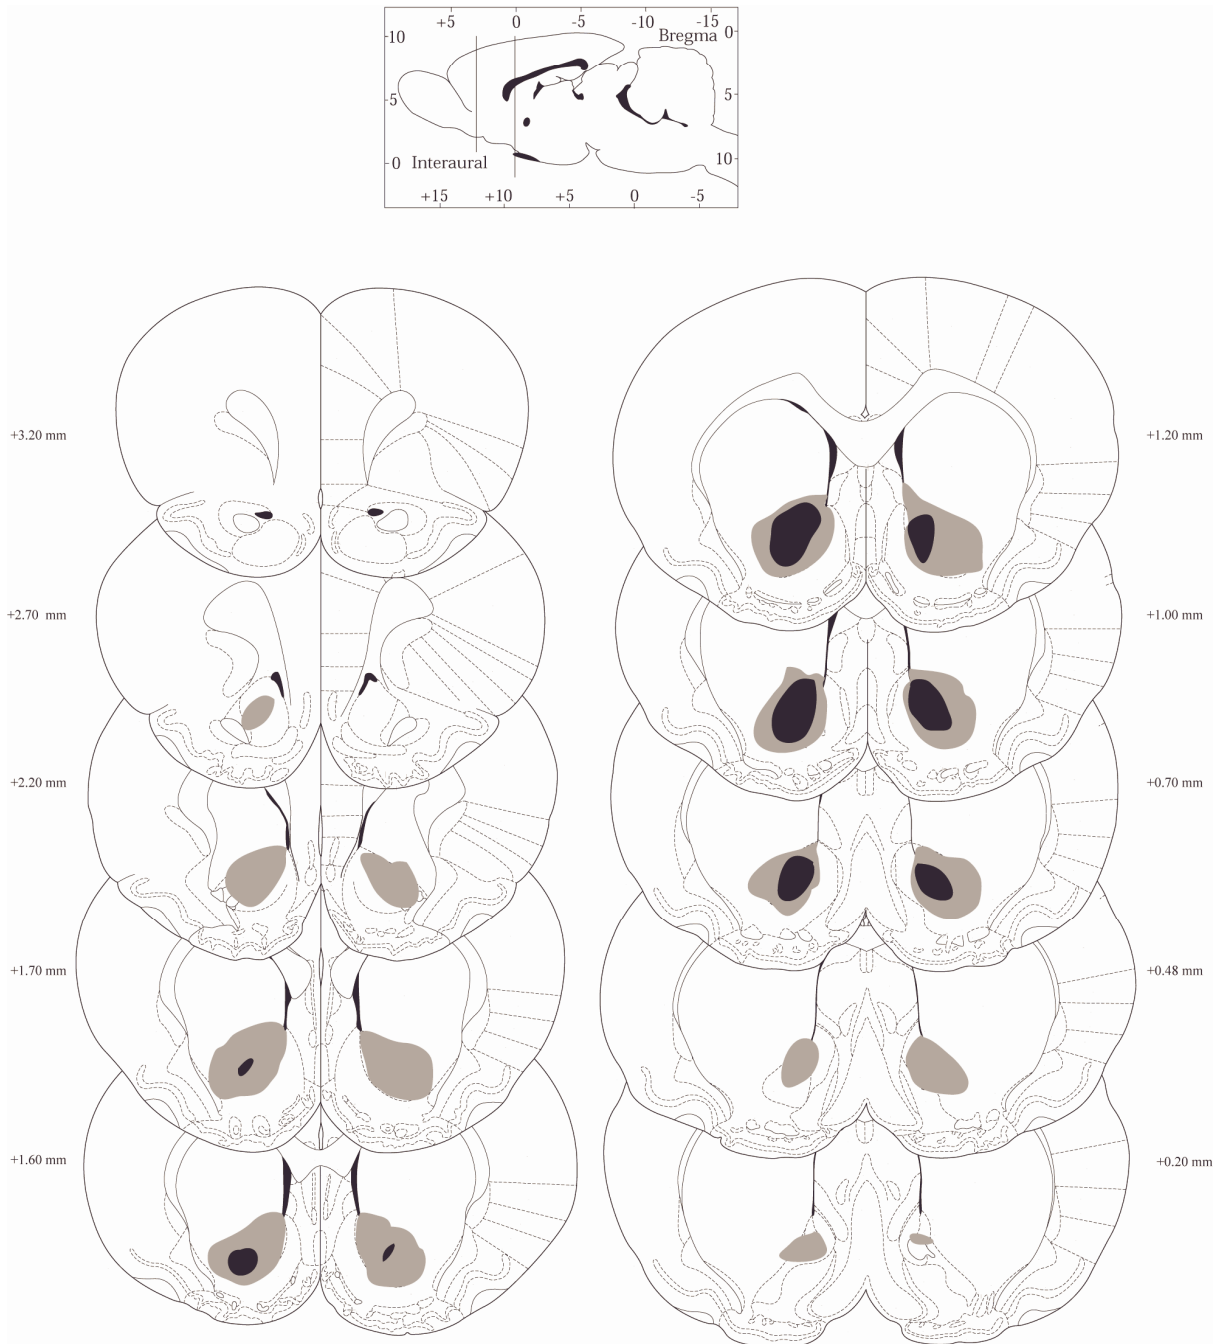


Figure 89. Lesions of the AcbC (subjects J1, J11, J12, J13, J15, J16, J17, J18, J21, J24). Grey shading indicates the extent of the largest area of neuronal loss, and black the smallest. Diagrams are taken from Paxinos & Watson (1998).

Body mass

The core-lesioned group appeared lighter than the control group, and this observation was confirmed. The groups did not differ in body mass at the time of surgery (means \pm SEMs: sham 356.0 ± 5.5 g, core 357.9 ± 8.0 g; one-way ANOVA: $F < 1$, NS). However, the core group were underweight thereafter, at around 90% of the mass of the control group — when feeding freely 6–13 days after surgery (sham 407.0 ± 5.6 g, core 359.2 ± 11.9 g; $F_{1,18} = 13.253$, $p = .002$), at the completion of the first delay-omission test (sham 371.8 ± 7.0 g, core 336.8 ± 7.9 g; $F_{1,18} = 10.941$, $p = .004$), and after all delayed reinforcement testing, at the start of the food consumption tests (sham 367.8 ± 7.1 g, core 324.4 ± 10.3 g; $F_{1,18} = 12.055$, $p = .003$).

Pre-operative acquisition

The groups remained matched after histological selection: there were no differences in the pre-operative pattern of choice (Figure 90A, p. 212). Choice ratios from the last 3 pre-operative sessions were analysed using the model group \times (delay \times S). While there was a highly significant effect of delay ($F_{3,124,56,234} = 20.542$, $\tilde{\epsilon} = .781$, $p < .001$), there was no effect of group, and no group \times delay interaction ($F_s < 1$, NS).

Baseline post-operative performance

Food was reliably collected and consumed, with the exception of a single occasion on which one core-lesioned subject (J13) was discovered to have left ~9 pellets uneaten at the end of the session.

Choice. The core-lesioned group were dramatically impaired in their ability to choose the large, delayed reward (Figure 90B, p. 212). An analysis of choice ratios from the 7 baseline sessions revealed a significant main effect of group ($F_{1,18} = 13.859$, $p = .002$) and a group \times delay interaction ($F_{4,72} = 2.964$, $p = .025$) in addition to a main effect of delay ($F_{4,72} = 37.28$, $p < .001$). However, subgroup analyses showed that both groups still exhibited a within-session shift in preference from the large to the small reward (sham group, effect of delay: $F_{4,36} = 23.668$, $p < .001$; core group: $F_{4,36} = 14.57$, $p < .001$).

As Figure 90B shows, the variance in the core group was substantially less than that in the sham group; the core group's preference for the immediate, smaller reinforcer was very consistent. As this heterogeneity of variance affected the ANOVA in which the two groups were compared (though not those analyses considering each group separately), non-parametric analyses were also conducted. Mann-Whitney U tests confirmed that the core group chose the delayed reinforcer significantly less often than shams at every single delay ($p < .023$ in each case). Surprisingly, the core group chose the large reinforcer less often than the small reinforcer at zero delay (comparison to 50%, $t_9 = -5.147$, $p = .001$).

To confirm that this change reflected a change in the performance of the core group, and not of the shams, choice ratios from the last 3 pre-operative sessions were compared with those from the first 3 post-operative sessions using the model group₂ \times (pre/post₂ \times delay₅ \times S). This revealed a significant pre/post \times group interaction ($F_{1,18} = 10.302$, $p = .005$). Separate analyses of the core and sham groups showed that the choice behaviour of the sham group did not alter following surgery ($F_{1,9} = 3.199$, $p = .107$), while that of the core group did ($F_{1,9} = 7.437$, $p = .023$).

Although Figure 90B suggests that the core-lesioned group exhibited a slightly reduced within-session shift in preference, because they rapidly approached reached a 'floor' at which the delayed reinforcer was seldom chosen, the rapidity of this shift (as assessed by the regression slope measure) did not differ between groups and did not alter across the post-operative sessions (Figure 90C). Analysis of the slope measures using the model group₂ \times (session₇ \times S) revealed a non-significant trend towards less steep (less negative/numerically greater) slopes in the core group (effect of group: $F_{1,18} = 3.624$, $p = .073$), with no effect of session and no interaction ($F_s < 1$, NS).

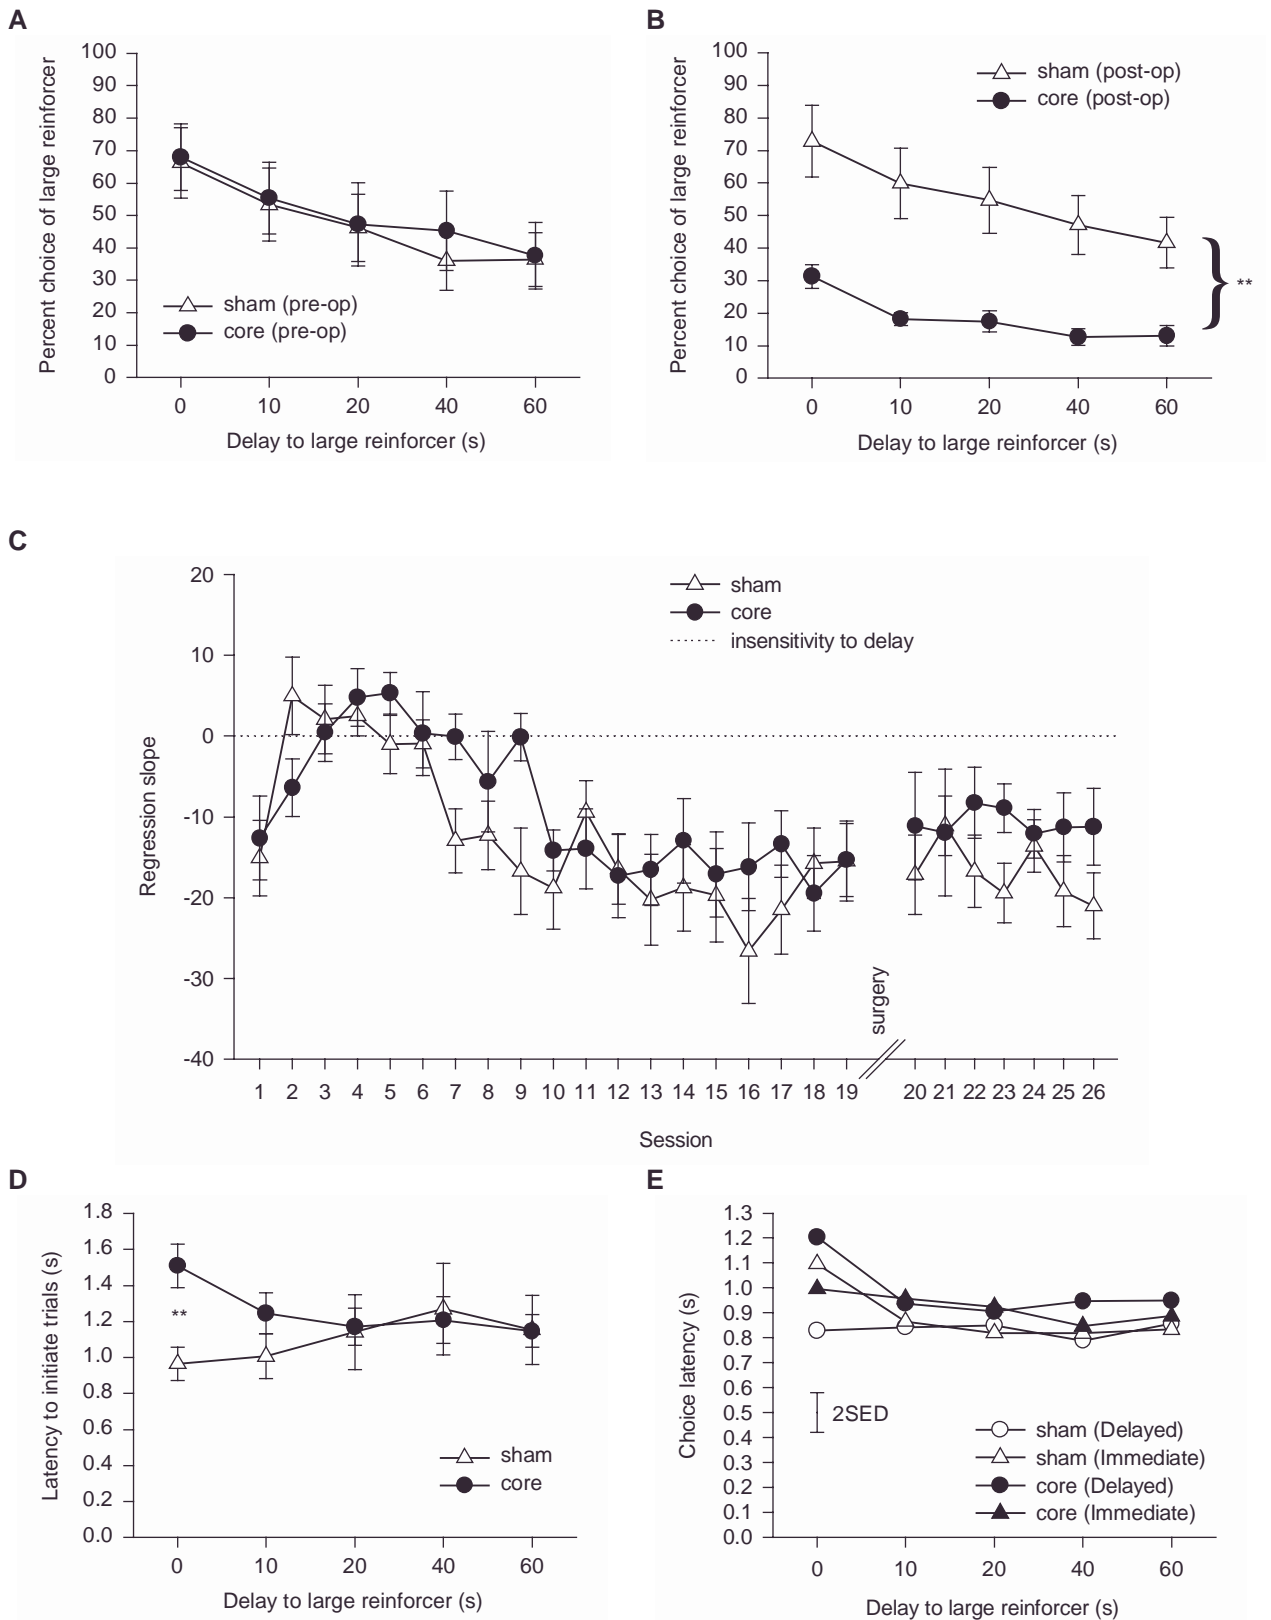


Figure 90. Performance of rats with lesions of the AcbC on the delayed-reinforcement choice task: baseline sessions. **A:** Pre-operative performance — data from the last 3 sessions preceding surgery. **B:** Post-operative performance — data from the first 7 sessions following surgery. Core-lesioned rats were significantly impaired in their ability to choose the larger, delayed reward. **C:** Slope measures before and after surgery. **D:** Latencies to initiate trials. **E:** Choice latencies. (2SED, twice the standard error of the difference for the three-way interaction; ** $p < .01$.)

Omissions. Responding was reliable, with all animals regularly sampling both levers, and the two groups did not differ in the number of omissions made. An analysis of the percentage of trials on which an omission occurred, across all delays, revealed no effect of group ($F_{1,18} = 2.665$, NS).

Initiation latency. The core-lesioned subjects were slower to initiate trials at zero delay (Figure 90D). Analysis of initiation latencies revealed a group \times delay interaction ($F_{2,548,45.868} = 5.274$, $\tilde{\epsilon} = .637$, $p = .005$). As Figure 90D suggests, this was due to slower initiation by the core group at zero delay (one-way ANOVA: $F_{1,18} = 12.903$, $p = .002$); between-group differences at non-zero delays were not significant ($p > .177$).

Choice latency. There was a complex but small difference between groups in choice latency (Figure 90E). Analysis of choice latencies using the model group₂ \times (response₂ \times delay₅ \times S) revealed a group \times response \times delay interaction ($F_{2,779,38.905} = 3.436$, $\tilde{\epsilon} = .695$, $p = .029$) in addition to a main effect of delay ($F_{1,636,22.897} = 4.344$, $\tilde{\epsilon} = .409$, $p = .032$). Analyses of the two group separately showed that there was a response \times delay interaction in the sham group ($F_{4,24} = 3.624$, $p = .019$), probably due to slower responding on the Immediate lever at zero delay (though the latency difference between the two levers at zero delay was not significant in its own right by *post hoc* testing; $F_{1,8} = 3.332$, $p = .105$), though this interaction was not significant in the core group ($F < 1$, NS).

Collection latency. The two groups did not differ in the speed with which they collected the rewards. An analysis of collection latencies using the model group₂ \times (response₂ \times delay₅ \times S) revealed no significant terms (maximum F was for the three-way interaction: $F_{2,987,41.82} = 2.504$, $\tilde{\epsilon} = .747$, $p = .072$).

Nosepoking during the delay. There were no differences between the two groups in the rate of nose-poking in the food alcove during the delay. An analysis using the model group₂ \times (delay₄ \times S) revealed no significant terms (maximum $F_{1,15} = 2.101$, NS).

Effect of omitting all delays (ABAB design)

Both groups remained sensitive to the delay. Removing the delays in alternating sessions increased both the sham- and core-lesioned groups' preference for the larger reward (Figure 91A, p. 214). Analysis of choice ratios using the model group₂ \times ({Delays versus No Delays}₂ \times trial block₅ \times S) revealed a highly significant interaction between the Delay/No Delay factor and the trial block ($F_{1,802,32.444} = 16.391$, $\tilde{\epsilon} = .451$, $p < .001$) in addition to main effects of group ($F_{1,18} = 8.238$, $p = .01$), the Delay/No Delay factor ($F_{1,18} = 15.622$, $p = .001$), and trial block ($F_{2,989,53.802} = 15.996$, $\tilde{\epsilon} = .747$, $p < .001$). The group \times {Delay/No Delay} interaction escaped significance ($F_{1,18} = 4.182$, $p = .056$), as did the three-way interaction ($F < 1$, NS). Heterogeneity of variance was not significant.

Confirming this statistical picture, a {Delay/No Delay} \times trial block interaction was detectable for both the sham group ($F_{4,36} = 13.768$, $p < .001$) and the core group ($F_{1,407,12.659} = 5.717$, $\tilde{\epsilon} = .352$, $p = .025$) when analysed separately.

The core group's preference for the larger reward remained significantly below that of the sham group in the No Delay condition (main effect of group in the No Delay condition: $F_{1,18} = 9.422$, $p = .007$).

Effect of omitting all delays (AAABBB design)

A further delay-omission test was conducted using three consecutive delay or no-delay sessions (AAABBB design). Although this more prolonged experience with the No Delay condition succeeded in increasing the core group's preference for the larger reward, the basic pattern remained the same as for the previous delay-omission test (Figure 91B).

Analysis identical to that for the previous test again detected a highly significant {Delay/No Delay} \times trial block interaction, and main effects of the {Delay/No Delay factor} and of trial block; however, in

this test, surprisingly, no group differences were significant (group: $F_{1,18} = 2.97$, $p = .102$; other terms involving group: $F < 1$, NS). Subgroup analyses demonstrated significant {Delay/No Delay} \times trial block interactions in both the sham and the core groups. In this test, however, the difference between the sham and core groups in the No Delay condition was not significant ($F_{1,18} = 2.567$, $p = .127$).




Figure 91. Effect of removing all delays on the performance of sham- and core-lesioned rats. **A:** Effect of omitting all delays in alternating sessions (ABAB design). **B:** Effect of omitting all delays with three consecutive sessions in each condition (AAABBB design). (2SED, twice the standard error of the difference for the three-way interaction. Note that this error term is *not* appropriate for the simple between-group comparison.)

Prolonged training without delays, and subsequent reintroduction of delays

Despite the lack of a statistical difference between the groups in the final delay omission test (Figure 91B), the core group's absolute level of preference for the larger reward was not as high as that of the sham group. In an attempt to equalize the groups, further training was given with no delays present (see *Methods*). Data from the last of these sessions are shown in Figure 92A. Though the two groups were not equalized by this training, all tendency to exhibit a within-session shift in preference was removed (Figure 92A). Subsequent reintroduction of delays caused preference for the larger reinforcer to collapse; Figure 92(B–D) shows consecutive blocks of 3 sessions. As the sessions proceeded, the core group's preference for the delayed reinforcer declined first at long delays, and then at progressively shorter delays.

Nevertheless, as clear preference for the large reinforcer had not been re-established in the core group as a whole, one further analysis was conducted. From the last day of no-delay training (session 42; Figure 92A), those rats were selected that met a criterion of $\geq 90\%$ choice of the large reinforcer in every trial block. This selection eliminated 3 rats from the sham group, leaving 7, and eliminated 5 rats from the core group, leaving 5 (Table 20). Having selected those rats that clearly discriminated between the two reinforcers and were not in the least biased away from the large-reinforcer lever as a result of their experience with delays, Figure 92 was replotted; the results are shown in Figure 93. It can be seen clearly that even those core-lesioned rats that exhibited a strong preference for the large reinforcer when it was delivered immediately (Figure 93A) were extremely intolerant of delay compared to the sham group (Figure 93D). The fact that these core-lesioned rats strongly preferred the large reinforcer in a task when no delays were present at all, but that their preference for the large reinforcer at zero delay declined when delays were reintroduced (compare Figure 93A and Figure 93D at zero delay) suggests that the severe deficit in the efficacy of delayed reward affected responding at non-zero delays, and then generalized to affect their

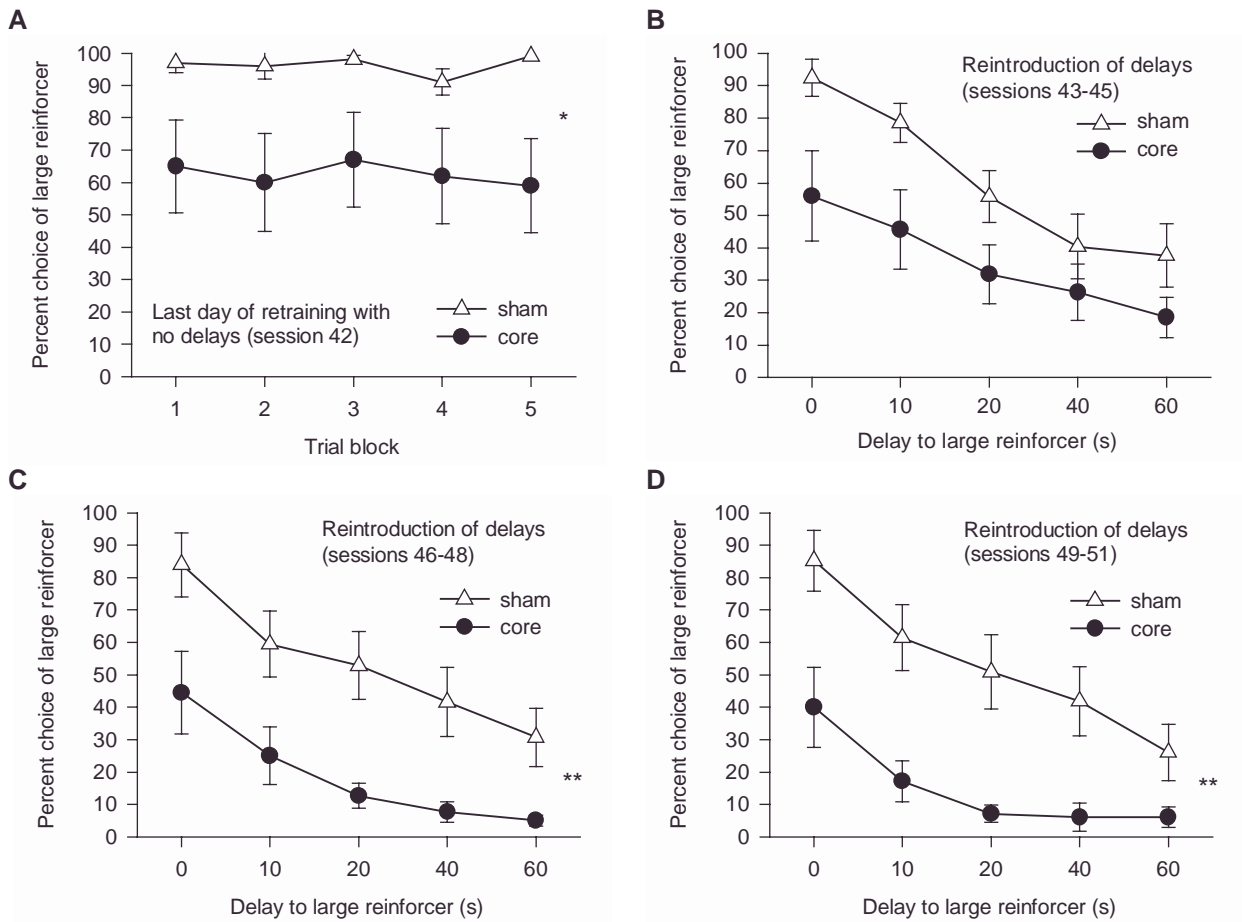


Figure 92. A: Preference following extended training in the absence of any delays (full data set shown in Table 20). B–D: Performance over consecutive blocks of sessions upon the reintroduction of delays. As these data exhibit significant heterogeneity of variance, the highly conservative correction of Box (1954) was applied (see Howell, 1997, pp. 322/457/464); * $p < .05$, ** $p < .01$ for the corrected between-group difference.

Table 20. Performance of sham-operated and core-lesioned subjects on the final day of extended training in the absence of delays (session 42, Figure 92A and Figure 93A). The percentage of trials on which the large reinforcer was chosen is shown, for each of the five blocks of ten choice trials. All sham-operated controls and the majority of AcbC-lesioned rats showed a preference for the large reinforcer (>50%) in all trial blocks. Rats that met the more stringent criterion of $\geq 90\%$ choice of the large reinforcer in every trial block were used for a further analysis (Figure 93). No omissions were made in this session.

Rat	J2	J3	J4	J7	J8	J9	J10	J14	J20	J22	J1	J11	J12	J13	J15	J16	J17	J18	J21	J24
Group	sham	sham	sham	sham	sham	sham	sham	sham	sham	sham	core	core	core	core	core	core	core	core	core	core
Trial block 1	100	100	100	100	100	100	100	70	100	100	90	0	100	100	0	10	100	100	100	50
Trial block 2	100	100	100	100	100	100	100	60	100	100	90	0	100	100	0	0	100	90	100	20
Trial block 3	100	100	100	100	100	90	100	90	100	100	90	0	100	100	0	0	100	90	100	90
Trial block 4	70	100	100	100	70	100	90	80	100	100	100	0	100	90	0	0	100	90	100	40
Trial block 5	100	100	100	100	90	100	100	100	100	100	100	0	100	70	0	0	100	100	90	30
>50% throughout?	√	√	√	√	√	√	√	√	√	√	√	×	√	√	×	×	√	√	√	×
$\geq 90\%$ throughout?	×	√	√	√	×	√	√	×	√	√	√	×	√	×	×	×	√	√	√	×

preference even in the zero-delay condition. This may explain why the core group demonstrated a deficit in responding for the large reinforcer even at zero delay during baseline testing sessions (Figure 90B, p. 212).

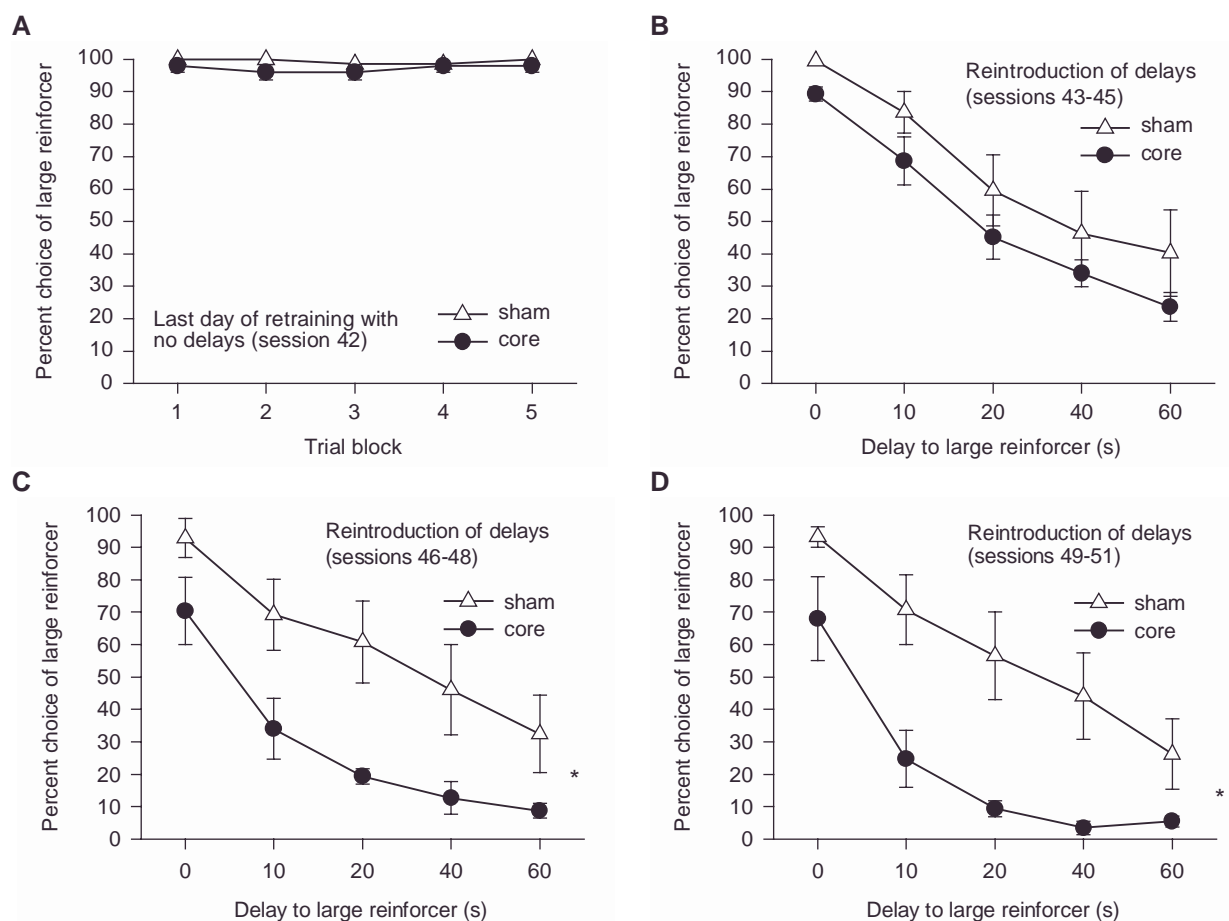


Figure 93. This figure is identical in form to Figure 92, but only includes data from those rats selected on the basis of a criterion of $\geq 90\%$ preference for the large reinforcer on the last day of training with no delays (see Table 20). The groups were therefore matched in panel A. In panels B–D, upon reintroduction of the delays, preference for the large reinforcer collapsed in the core group. As these data exhibit significant heterogeneity of variance, the highly conservative correction of Box (1954) was applied (see Howell, 1997, pp. 322/457/464); * $p < .05$ for the corrected between-group difference.

Locomotor activity in a novel environment

Core-lesioned subjects were hyperactive, and slower to habituate to the novel environment of the locomotor testing apparatus (Figure 25). Following square-root transformation, analysis of the total number of infrared beam interruptions using the model $\text{group}_2 \times (\text{bin}_{12} \times S)$ revealed an effect of bin ($F_{7,777,139,994} = 12.079$, $\tilde{\epsilon} = .707$, $p < .001$), reflecting habituation, but also an effect of group ($F_{1,18} = 12.057$, $p = .003$), and a group \times bin interaction ($F_{7,777,139,994} = 2.279$, $\tilde{\epsilon} = .707$, $p = .027$).

Food consumption tests

The core-lesioned subjects ate more slowly than the sham-operated controls, at least when consuming the chow used as their maintenance diet; differences in food consumption were not significant for the sucrose pellets used in the delayed reinforcement task.

Mass of chow consumed in 30 min. There was a small but significant difference in the amount of chow consumed: the core group ate less. The mean \pm SEM amounts consumed were 8.0 ± 0.4 g (sham) and 6.5 ± 0.5 g (core); one-way ANOVA demonstrated these to be significantly different ($F_{1,18} = 5.777$, $p = .027$).

Time to consume 2.25 g chow. The core group ate the fixed amount of chow more slowly (501 ± 39 s) than the shams (375 ± 6 s). Inhomogeneity of variance necessitated a nonparametric test; the difference between the two groups was significant by a Mann-Whitney U test ($p = .005$).




Figure 94. Locomotor activity in a novel environment (120-min session scored in 10-min bins). The core group were hyperactive and habituated more slowly (** $p < .01$, main effect of group; # $p < .05$, group \times bin interaction).

Mass of sucrose pellets consumed in 30 min. The core group ate less (9.1 ± 0.5 g) than the shams (11.4 ± 0.9 g); however, this difference was not significant (inhomogeneity of variance necessitated a non-parametric test; Mann-Whitney U test, $p = .063$).

Time to consume 50 sucrose pellets (2.25 g). Though the core group ate the fixed mass of sucrose pellets more slowly (250 ± 30 s) than the shams (199 ± 20 s), this difference was not significant ($F_{1,18} = 1.964$, $p = .178$).

Summary

Lesions of the AcbC induced a profound and lasting deficit in subjects' preference for the large reward when it was delayed. Subjects remained sensitive to removal of the delay and discriminated the two reinforcers. In baseline testing sessions, AcbC-lesioned subjects also failed to choose the large reward as often as shams when it was not delayed; however, prolonged training in the absence of delays re-established preference for the large reinforcer in a majority of lesioned subjects, and these subjects remained hypersensitive to the effects of reintroducing the delays subsequently. In addition, AcbC-lesioned rats were hyperactive, ate less of the food used as their maintenance diet (but showed normal consumption of the reinforcer used in the task), and were approximately 10% lighter than shams.

EXPERIMENT 4. EFFECTS OF INTRA-ACCUMBENS AMPHETAMINE ON CHOICE OF SIGNALLED AND UNSIGNALLED DELAYED REINFORCEMENT

Methods

Twenty-four naïve subjects were trained to press levers for sucrose pellets as before, and to nosepoke in order to initiate discrete-trial presentations of the levers, before being trained a variant of the delayed-reinforcement task adapted for intracranial infusions.

Abbreviated delayed-reinforcement task for intracranial infusions

This task was identical to the delayed-reinforcement choice procedure in Chapter 6, except that only three blocks of trials were used (each comprising two forced and ten free-choice trials), with a descending order of delays. This modification was made in an attempt to ensure that high Acb levels of drug coincided with responding at non-zero delays. The delays used were 60 s, 20 s and 0 s (in order). Trials began every 100 s, as before, for a total session length of 60 min.

Half the subjects were trained in the Cue condition ($n = 12$), and half in the No Cue condition ($n = 12$).

A stability criterion was defined as follows: after excluding single-lever trials, choice ratios (delayed lever responses ÷ total responses) were calculated for each rat using the summed responses for three consecutive sessions, and subjected to ANOVA with delay as a within-subjects factor. When the effect of delay was significant at the $\alpha = .01$ level, the rats were considered to have criterion performance from the first session of the three. (Note that this criterion is not exactly comparable to that used in Chapter 6, in light of the different group sizes used.) Following the triplet of sessions in which the criterion was attained, subjects were given 5 more baseline sessions on the task before surgery.

All subjects then received cannulae aimed at the Acb (see *Methods*). Following recovery, they were retrained on the basic task for 3 sessions, and given a single preliminary infusion of saline to accustom them to the infusion procedure (as described in Chapter 3, p. 77). The preliminary infusion was given in the testing room containing the operant chambers, but the subjects were returned to their home cages following infusion.

Intra-accumbens amphetamine. Four doses of *d*-amphetamine sulphate (0, 3, 10, 20 μg) were given in a volume of 1 μl bilaterally in a digram-balanced Latin square, immediately before each test session. The infusion procedure was described in detail in Chapter 3 (p. 77). The Latin square was then repeated, in order to accumulate data from two sessions per dose per rat.

Results

Regrettably, three rats in the Cue group (N17, N18, N19) died post-operatively, as did one rat (N5) in the No Cue group. One other rat in the No Cue group (N6) died during behavioural testing, and its data were discarded.

Histology

On the whole, the cannula tips were located more ventrally than in previous experiments; they were positioned predominantly in the inferior shell, or at the core-shell boundary. Two rats with tip locations in the ventral pallidum (subjects N4, N24) were excluded, leaving 9 subjects in the No Cue group (rats N1, N2, N3, N7, N8, N9, N10, N11, N12) and 8 in the Cue group (rats N13, N14, N15, N16, N20, N21, N22, N23). Representative photomicrographs of Acb cannulae tracks and injector tip locations were shown in Chapter 3 (p. 82); schematics of the tip locations in the two groups are shown in Figure 95.

Schematic of cannula locations



Figure 95. Location of the tips of injection cannulae within the Acb. *Black crosses* indicate subjects in the No Cue group (subjects N1, N2, N3, N7, N8, N9, N10, N11, N12). *Red triangles* indicate subjects in the Cue group (subjects N13, N14, N15, N16, N20, N21, N22, N23). Diagrams are taken from the atlas of Paxinos & Watson (1998).

Acquisition and baseline performance

The Cue group acquired sensitivity to the programmed delay earlier than the No Cue group. The Cue group first met the $\alpha = .01$ delay-sensitivity criterion for sessions 7–9 and were operated following session 14, while the No Cue group met the criterion for sessions 24–26, and were operated following session 31.

The earlier acquisition in the Cue group was not apparent from an analysis of regression slopes during acquisition (Figure 96A, p. 221). Analysis of regression slope measures for the first 14 sessions (when both groups were in the pre-operative acquisition phase) using the model $\text{group} \times (\text{session} \times S)$ revealed an effect of session ($F_{10,849,162.73} = 4.144$, $\tilde{\epsilon} = .835$, $p < .001$), but no group effect and no $\text{group} \times \text{session}$ interaction ($F_s < 1$, NS).

However, consideration of choice behaviour did establish that the two groups showed different levels of performance at an equivalent time in the course of acquisition (Figure 96B); at this time, the absolute, session-wide level of preference for the delayed reinforcer was greater in the Cue group than the No Cue group even though the rapidity of the within-session shift in preference did not differ substantially. Analysis of choice ratios from sessions 12–14 in each group (at which time the Cue group had reached criterion but the No Cue group had not) revealed a significant difference in choice behaviour: statistically, there was a main effect of group ($F_{1,15} = 8.025$, $p = .013$) as well as of delay ($F_{1,566,23.496} = 10.327$, $\tilde{\epsilon} = .783$, $p = .001$), but no interaction ($F < 1$, NS).

This early difference between the groups disappeared as a result of further training of the No Cue group (Figure 96C). Comparison of choice ratios from the last 3 pre-operative sessions in each group (namely sessions 12–14 in the Cue group and sessions 29–31 in the No Cue group) yielded no group differences (delay: $F_{1,603,24.039} = 13.94$, $\tilde{\epsilon} = .801$, $p < .001$; group and $\text{group} \times \text{delay}$: $F_s < 1$, NS).

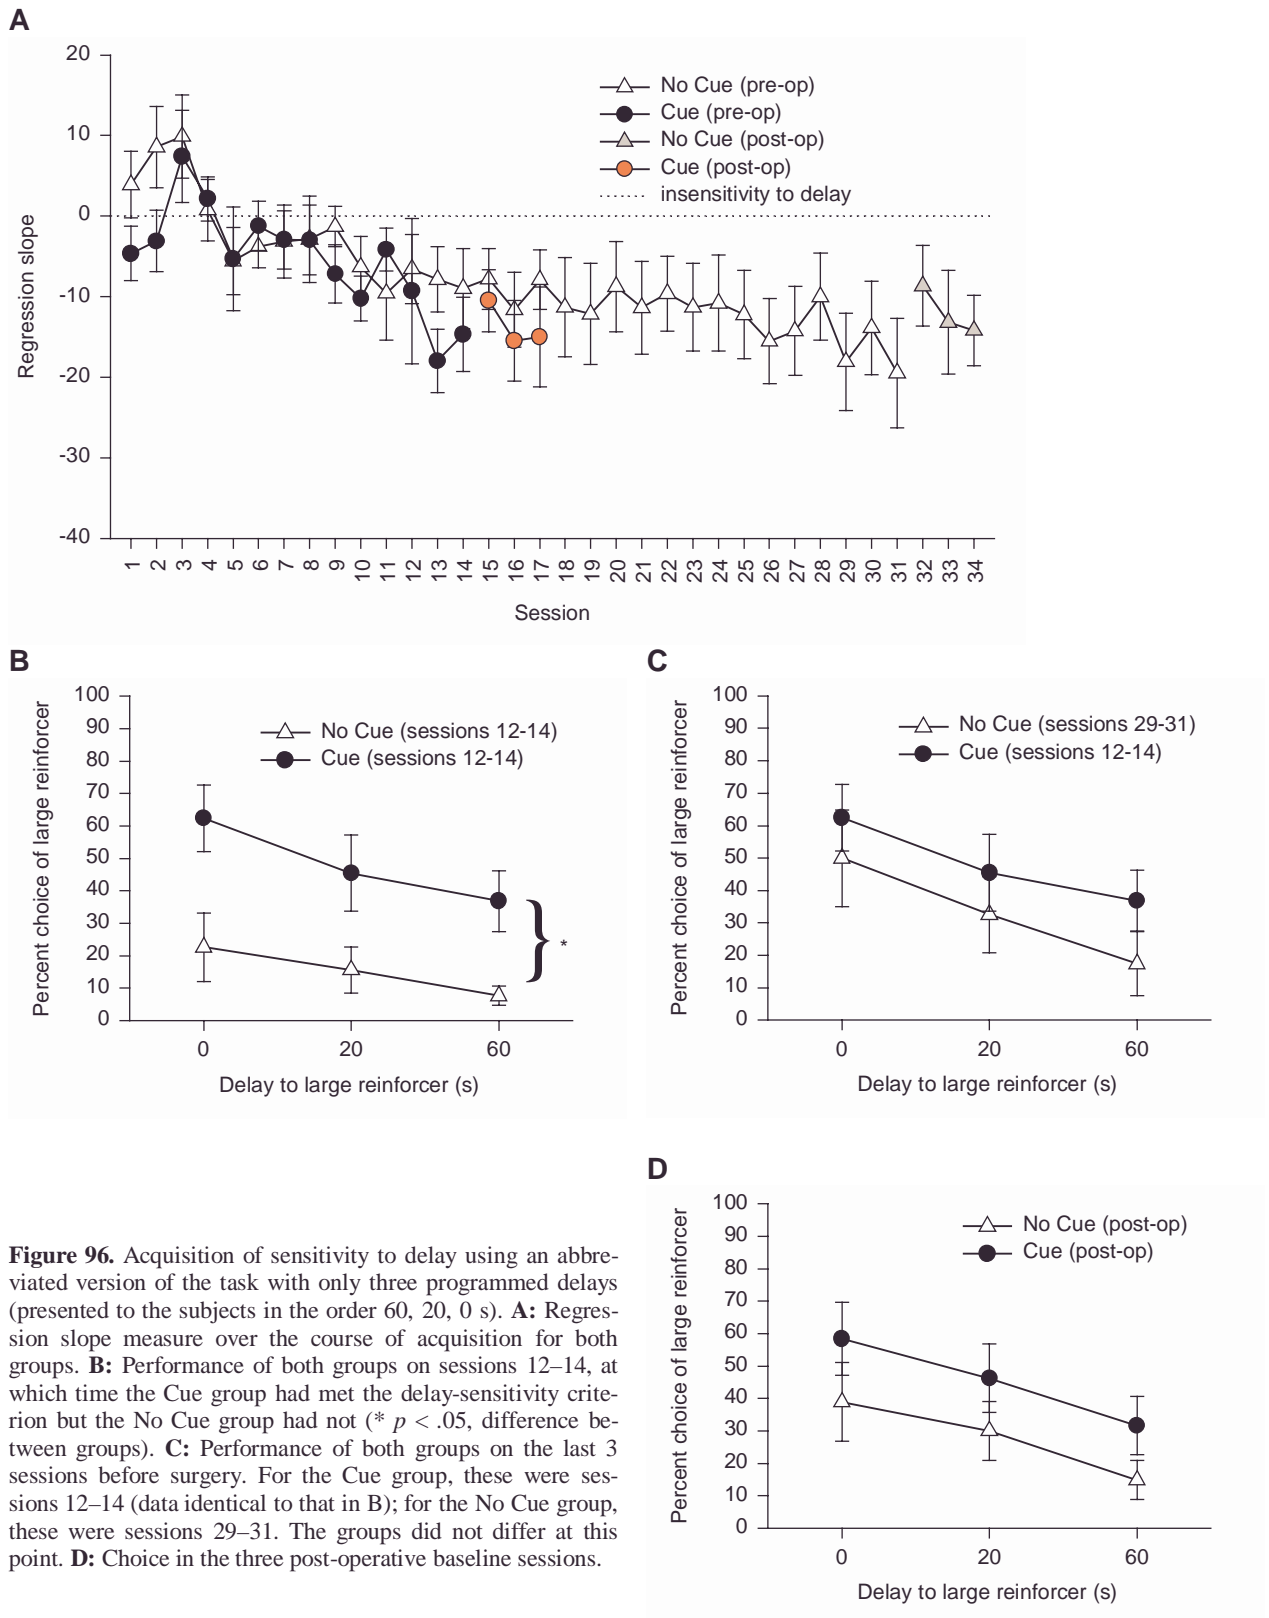
Re-establishment of baseline performance following surgery

Group differences did not re-emerge following surgery, either in choice ratio analysis (group: $F_{1,15} = 1.798$, NS; $\text{group} \times \text{delay}$: $F < 1$, NS; delay: $F_{2,30} = 18.592$, $p < .001$) or analysis of the regression slope measure, which was stable post-operatively (analysed using session and group as factors: all $F_s < 1$, NS).

Effects of intra-accumbens amphetamine on choice

Some doses of amphetamine decreased preference for the large, delayed reinforcer (Figure 97), particularly at the 20-s delay, but a cue-dependent effect was not found. Analysis of choice ratios using the model $\text{group}_2 \times (\text{dose}_4 \times \text{delay}_3 \times S)$ demonstrated main effects of dose ($F_{3,45} = 4.338$, $p = .009$) and delay ($F_{1,377,20.66} = 37.738$, $\tilde{\epsilon} = .689$, $p < .001$). The $\text{dose} \times \text{delay}$ interaction just escaped significance ($F_{6,90} = 2.169$, $p = .053$). No other term was significant ($F_s \leq 1.068$, NS).

Surprisingly, pairwise comparisons established that the 3 μg dose and the 20 μg dose significantly reduced preference for the large, delayed reinforcer ($p = .024$ and $.037$ respectively) relative to vehicle, while 10 μg had no effect ($p = .591$). The effects of 3 μg and 20 μg did not differ from each other ($p = .226$).



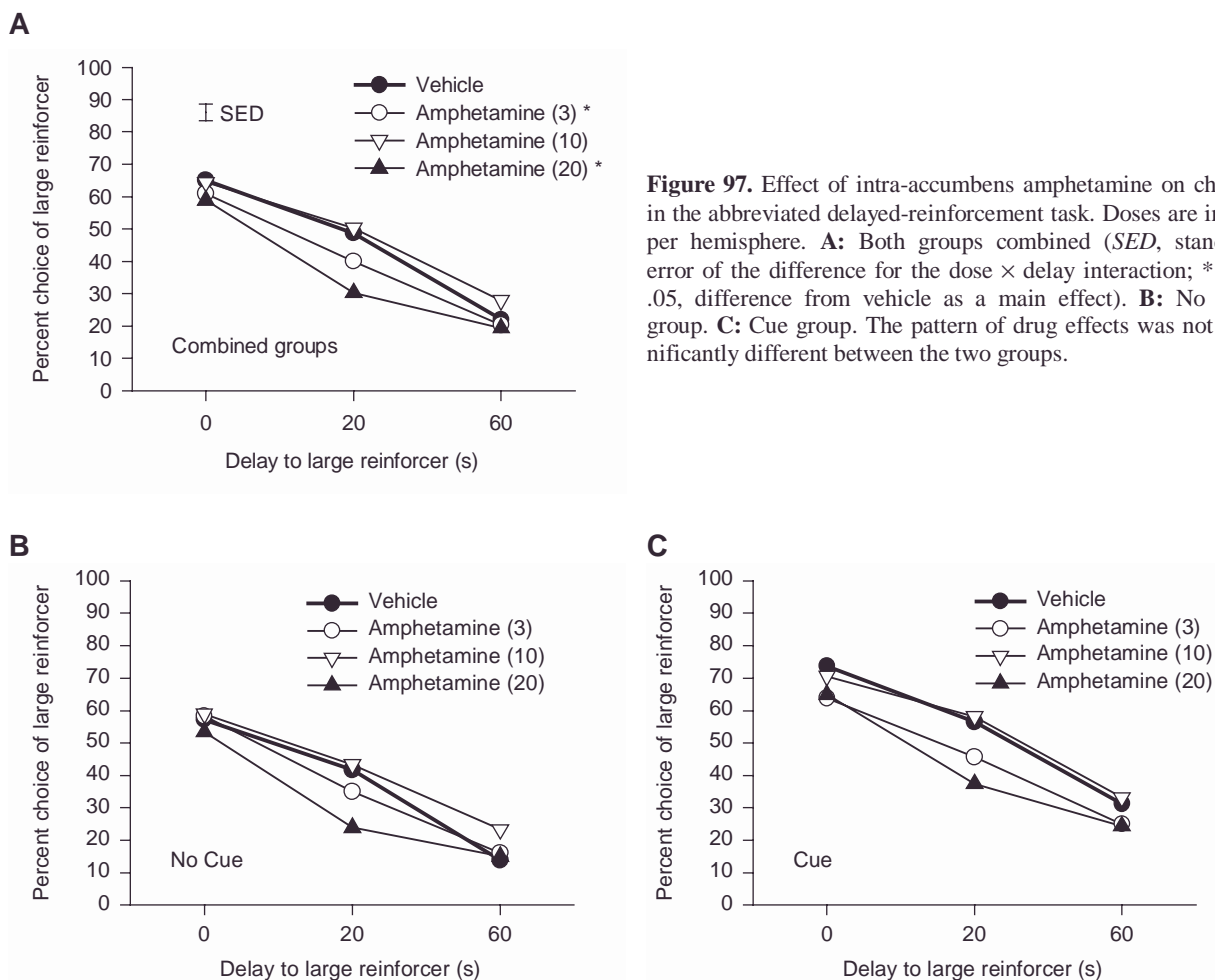


Figure 97. Effect of intra-accumbens amphetamine on choice in the abbreviated delayed-reinforcement task. Doses are in μg per hemisphere. **A:** Both groups combined (*SED*, standard error of the difference for the dose \times delay interaction; * $p < .05$, difference from vehicle as a main effect). **B:** No Cue group. **C:** Cue group. The pattern of drug effects was not significantly different between the two groups.

Effects of intra-accumbens amphetamine on latencies and nosepoking behaviour

Initiation latency. Though initiation latencies increased with delay — despite delays decreasing as the session progressed, so that trials were initiated faster at the end of the session — amphetamine did not affect the latency. Initiation latencies were analysed using the model group \times (dose \times delay \times S). There was a main effect of delay ($F_{1.685,25.272} = 8.449$, $\tilde{\epsilon} = .842$, $p = .002$) but no other term was significant (dose \times delay \times group: $F_{6,90} = 1.658$, NS; other $F_s < 1$, NS). The mean initiation latencies (in seconds) were 0.998 ± 0.089 (0 s), 1.019 ± 0.087 (20 s), and 1.235 ± 0.111 (60 s).

Choice latency. Amphetamine did not affect the latency to choose a lever. There were insufficient data to allow a full model to be used, so they were analysed as group \times (dose \times response \times S). This revealed no significant effect of any term (group: $F_{1,14} = 3.531$, $p = .081$, with a slight tendency for faster responding in the Cue group; other terms: $F \leq 2.114$, $p \geq .168$).

Collection latency. Subjects collected the immediate reward faster, and there was a non-significant tendency for amphetamine to slow collection of the large reward dose-dependently. Again, there were insufficient data for a full model, so group \times (dose \times response \times S) was used. This revealed a near-significant dose \times response interaction ($F_{1.65,23.106} = 2.904$, $\tilde{\epsilon} = .55$, $p = .083$) in addition to a main effect of response ($F_{1,14} = 39.006$, $p < .001$).

Nosepoking during the delay. Subjects nosepoked for a greater proportion of the 20-s delay (15%) than of the 60-s delay (12%), and it appeared that the highest dose of amphetamine reduced nosepoking (means across both delays for each dose: vehicle 13.8%, 3 μ g 14.9%, 10 μ g 14.0%, 20 μ g 9.8%). Although the Cue group did nosepoke for more of the delay than the No Cue group (15.2% versus 11.0% respectively), as in Chapter 6, this difference was not significant. Using the proportion of the delay spent nosepoking as the dependent measure, an analysis using the model $\text{group}_2 \times (\text{dose}_4 \times \text{delay}_2 \times S)$ was conducted. This revealed main effects of dose ($F_{3,24} = 4.757, p = .01$) and of delay ($F_{1,8} = 11.929, p = .009$), with no other significant terms ($F_s \leq 1.07, \text{NS}$). However, using Sidak-corrected pairwise comparisons, no single dose was found to be significantly different from any other in *post hoc* tests ($p > .12$).

Summary

An abbreviated version of the delayed reinforcement choice task was used for this experiment, with a descending order of delays. Intra-Acb amphetamine reduced subjects' preference for the large, delayed reward slightly, but not in a clear dose-dependent manner (with effects being observed at 3 μ g and 20 μ g, but not at 10 μ g). The effects of amphetamine were not demonstrably different in groups trained with and without a cue present during the delay. The 20- μ g dose of amphetamine also appeared to have slight effects to reduce nosepoking in the food alcove during the delay to reinforcement.

DISCUSSION

Lesions of the AcbC induced a profound, long-lasting deficit in the ability to choose a delayed reward; these rats responded reliably but made highly impulsive choices. In contrast, lesions of the mPFC induced a subtle deficit in the pattern of responding while lesions of the ACC had no effect on choice. These experiments represent the first use of focal excitotoxic lesions to study choice of delayed reinforcement, and used a technique of matching corresponding sham and lesioned groups for performance pre-operatively, ensuring high power to detect changes due to the lesions. Intra-accumbens amphetamine injections had somewhat inconsistent effects to reduce preference for the delayed reward, and this effect did not depend on whether the delay was bridged by a signal. The effects of each manipulation will first be discussed separately.

Effects of anterior cingulate cortex lesions

Lesions of the ACC had no effect on choice, establishing that the ACC is not required for rats to choose a delayed reinforcer. Moreover, ACC-lesioned rats remained equally sensitive to unexpected removal of the delays in this task, suggesting that their choices were no more inflexible or 'habitual' than those of shams.

This finding stands in apparent contrast to previous reports of motor impulsivity or disinhibited responding in ACC-lesioned rats. For example, such rats have been found to over-respond to unrewarded stimuli (Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c), and to respond prematurely in situations where they are required to wait (Muir *et al.*, 1996) (though the present lesions may be different from those of Muir *et al.*; see also Chapter 3, p. 119). However, such a dissociation is not in itself unexpected, as motor impulsivity and impulsive choice have been dissociated before ('execution' and 'outcome' impulsivity; Evenden, 1999b).

The ACC-lesioned rats were slower to collect the larger reward, the only behavioural effect of these lesions evident in this task. This deficit resembles very closely the increased latency of ACC-lesioned rats to approach the CS+ predictive of food observed in the autoshaping tasks used in Chapter 3 (and discussed there, p. 97). The slowing might reflect damage to the motor regions of the ACC (Dum & Strick, 1993), but in the present task other measures of response speed (trial initiation and choice latency) were not affected, suggesting perhaps that approach behaviour in ACC-lesioned rats was no longer enhanced by the expectation of a large reward.

The present results also provide a degree of further support for the hypothesis developed in Chapter 3 that the ACC is not critical for instrumental discrimination. Lesioned subjects in the present experiment discriminated between the two levers as well as control subjects did, despite the levers' being visually identical aside from their left/right position. This is in accordance with the view that the ACC is primarily important for the discrimination of similar Pavlovian conditioned stimuli on the basis of their association with reward. As discussed in Chapter 3 (p. 113), there is also evidence to suggest that the ACC may only play a critical role *early* in the learning of some tasks. It is of course possible that this applies to the delay-of-reinforcement task; the lack of a lesion effect does not preclude the involvement of the ACC in task acquisition. However, these results do suggest, despite findings of ACC abnormalities in disorders of impulsivity (e.g. Bush *et al.*, 1999), that ACC dysfunction is not an important contributor to impulsive choice.

Effects of medial prefrontal cortex lesions

Lesion of the mPFC 'flattened' the within-session shift from the large to the small reward; the mean preference for the large reward was *below* that of shams at zero delay, but *above* that of shams at the maximum delay. There is no obvious explanation for this effect within theories of choice of delayed reinforcement; it seems clear that the mPFC lesion resulted in some form of insensitivity to the contingencies or stimuli present in the task.

Contingency perception

Given that Balleine & Dickinson (1998a) demonstrated that lesions encompassing prelimbic cortex impaired rats' sensitivity to instrumental contingencies, it would be reasonable to suggest that a failure of contingency perception was responsible for performance of mPFC-lesioned rats in the present task. However, these rats were just as sensitive as controls to the unexpected removal of all delays; their responding was not inflexible, as might have been expected according to this account. The mPFC group were generally slower to respond on the levers, but this cannot easily be related to a specific cognitive deficit.

Timing ability

A more plausible interpretation is that mPFC lesions disrupted the control over behaviour by the passage of time in each session. There is strong evidence that normal rats learn a session-wide temporal discrimination in this task, and that this temporal discriminative stimulus comes to control responding, and in particular the tendency to shift from the large to the small reward as the session progresses (Chapter 6; Cardinal *et al.*, 2000b). Disruption of such temporal stimulus control might be expected to produce a flattening of the within-session shift of the kind seen.

Indeed, aspirative lesions of the mPFC have previously been shown to induce a general deficit in timing ability in rats (Dietrich & Allen, 1998); lesioned subjects showed a temporal discrimination function that was less steep than normal in the peak procedure, an operant task that assesses the ability to time a discriminative stimulus (Catania, 1970; Roberts, 1981). Indeed, 'temporal organization of behaviour' (albeit an ill-defined term) has been suggested to be a cardinal function of the prefrontal cortex (see e.g. Fuster, 1995). While disruption of timing behaviour on a shorter scale might in principle also affect choice behaviour in a delay-dependent manner (as discussed below, p. 229), there was no evidence for this in mPFC-lesioned subjects.

Effects of nucleus accumbens core lesions

Lesions of the AcbC induced a major deficit in subjects' ability to choose a delayed reward; lesioned subjects made truly impulsive choices. This was not due to an inflexible bias away from the lever producing the delayed reinforcer: AcbC-lesioned rats still chose the large reinforcer more frequently at zero delay than at other delays, and removal of the delays resulted in a rapid and significant increase in the rats' preference for the large reinforcer. Thus, the pattern of choice genuinely reflected a dramatically reduced preference for the large reinforcer when it was delayed, suggesting that delays reduced the effectiveness or value of rewards much more in AcbC-lesioned rats than in controls.

In the initial set of post-operative sessions, the AcbC-lesioned rats preferred the small reinforcer even at zero delay, avoiding the large reinforcer. Prolonged training in the absence of delays did not overcome the tendency to avoid the lever previously associated with delayed reinforcement in all lesioned subjects. Given the pre-operative performance of the same animals (i.e. equal to that of controls), this suggests that the post-operative experience of delayed reinforcement may have been highly aversive for AcbC-lesioned rats (or at least, much less preferable than immediate small reinforcement), inducing them to avoid that

lever permanently. However, the majority of core-lesioned subjects (6 out of 10) showed a consistent preference for the large reinforcer after prolonged training without delays (Table 20, p. 215). Even when sham and AcbC-lesioned subjects were selected who showed near-exclusive preference for the large reinforcer under these conditions, reintroduction of delays caused a dramatic and selective fall in preference for the large, delayed reinforcer in the AcbC-lesioned group (accompanied by a small decline in preference at zero delay; Figure 93, p. 216). These results suggest that the AcbC-lesioned rats' low preference for the large reinforcer at zero delay in the baseline post-operative sessions (Figure 90B, p. 212) was *not* due to a genuine preference for the small reward over the larger reward. Instead, it suggests that this result reflected the marked effects of the delays present later in the session (discussed further below, p. 228).

Primary motivational changes

AcbC-lesioned rats were underweight, and at least two possible contributing factors were observed: these rats exhibited locomotor hyperactivity and ate less of the chow used as their maintenance diet. These changes have been observed before following AcbC lesions (Parkinson, 1998). It is therefore possible that the lesioned rats' motivation to earn food was lower. However, it is unlikely that these changes contributed to their impulsive choice. First, there were no significant differences in the rate at which these subjects consumed the sucrose pellets used as the reinforcer in the task. Second, explicit manipulation of deprivation state has been shown not to affect choice on this task (Chapter 6; Cardinal *et al.*, 2000b). Third, performance of Acb-lesioned animals was not comparable in other respects to that of sated rats (Chapter 6; Cardinal *et al.*, 2000b); for example, they did not make more omissions than sham-operated controls.

Altered sensitivity to reinforcer magnitude or delay?

The core group showed at least some discrimination between the large and the small reinforcer. This is consistent with the observation that the expectancy of reward magnitude continues to have normal effects upon rats' reaction time following excitotoxic Acb lesions, with a smaller reaction time when a large reward is expected (Brown & Bowman, 1995) (though intra-Acb NMDA antagonists do impair this effect; Hauber *et al.*, 2000). A large proportion of the core group showed a preference, sometimes absolute, for the large reward when prolonged training was given with no delays. Five out of ten core-lesioned rats met a very stringent criterion for preference of the large reward under these conditions. These same rats were exquisitely sensitive to delays, preferring the large reinforcer much less than shams when it was delayed. Nevertheless, it remains a possibility that the other rats in the core group did not discriminate between the two reward magnitudes post-operatively, and that the history of delayed reinforcement on one lever permanently reduced their preference for that alternative.

It is also possible that the core group discriminated between the reinforcer magnitudes, but to a lesser extent than normal rats. In this scenario, core-lesioned rats still exhibit impulsive choice behaviourally — that much is clear — but because the perceived value of the large reinforcer is insufficient to overcome the normal effects of delay discounting. The multiplicative hyperbolic model of choice (see Ho *et al.*, 1999) postulates that the value of an immediate reinforcer of physical magnitude q is determined by the equation

$$V_{\text{immediate}} = \frac{q}{q + Q}, \text{ also expressed as } V_{\text{immediate}} = \frac{1}{1 + Q/q} \quad (1)$$

and that the value of this reinforcer when delayed by a time d is given by

$$V_{\text{delayed}} = \frac{V_{\text{immediate}}}{1 + K \cdot d} \quad (2)$$

where K and Q are ‘delay discounting’ and ‘magnitude discounting’ parameters that reflect intrinsic properties of the animal. In this theory, when one assesses an animal’s relative preference between two reinforcers of different magnitudes, one of which is delayed, changes in both K and Q may affect choice, as illustrated in Figure 98 and Figure 99. It can immediately be seen from these figures that both hypothetical kinds of manipulation can reduce preference for delayed rewards, inducing impulsive choice, though only one manipulation varies the effects of delay.




Figure 98. Hypothetical choice/delay curves for three individuals whose sensitivity to delay *per se* is identical, but across whom sensitivity to reinforcer magnitude varies. These curves were generated by assuming that the individual are offered an immediate reinforcer of magnitude 1, and a delayed reinforcer of magnitude 4. The absolute value to the animal of each reinforcer (V_1 and V_4) is calculated separately according to the hyperbolic discounting equations given in the text, and the relative preference is calculated as $V_4 / (V_1 + V_4)$. The delay sensitivity parameter K is identical in all three subjects, but the magnitude sensitivity parameter Q takes the values 0.01, 1, and 100. As $Q \rightarrow 0$, the animal becomes indifferent between the two reinforcers at zero delay; as $Q \rightarrow \infty$, relative valuation of reinforcers at zero delay approaches the relative proportion of their physical magnitudes (in this case, $4/(1+4)$ or 0.8).

Figure 99. Hypothetical choice/delay curves for three individuals whose sensitivity to reinforcer magnitude is identical, but across whom sensitivity to reinforcer delay varies. These curves were generated as for Figure 98, but Q is held constant (at 100) and K is varied (taking values of 0.1, 1, and 10, though the units are arbitrary).

Theoretically, a critical test of whether a given manipulation affects delay (K) or magnitude (Q) discounting, in this model, is to examine preference at zero delay, which manipulations of K cannot affect. Inspection of choice-by-delay plots (Figure 90 to Figure 93, pp. 212–216) suggests that lesions of the AcbC affected the perception of reinforcer magnitude, as preference for the large reinforcer at zero delay was not as high as that of shams. (A different interpretation is offered below.) Another, more direct test would be to obtain estimates of Q and K for each rat directly, and compare these across groups. In the present task, this might be attempted by assuming

$$\begin{aligned}
 \text{relative preference at delay } d &= \frac{V_{\text{large}}}{V_{\text{large}} + V_{\text{small}}} & (3) \\
 &= \frac{\frac{V_{\text{large-immediate}}}{1 + K \cdot d}}{\frac{V_{\text{large-immediate}}}{1 + K \cdot d} + V_{\text{small-immediate}}}
 \end{aligned}$$

This equation might be solved for physical reinforcer magnitudes (q) of 1 and 4 pellets and fitted to individual rats' data using non-linear programming techniques. However, this attempt is doomed to failure — not only because of the variability in rats' preferences, and by the poorly-constrained curve-fitting problem, but because it is clear that rats' preferences in the present task do not conform to this model. If choice ratios are interpreted as *relative preference* according to equations (1) and (3), a contradiction is apparent from Figure 92A (p. 215). Without delays, sham subjects' preferences approached 100% choice of the large reinforcer, whereas in the model, relative preference between a 1-pellet and a 4-pellet reinforcer cannot exceed 80%. The behavioural result comes as no surprise, for it is the well-known phenomenon of maximization on discrete-trial schedules (see Mackintosh, 1974, pp. 190–195).

Thus, behaviour on this task cannot be quantified according to the hyperbolic discounting model. A far more likely interpretation of the failure of core-lesioned rats to choose the large reinforcer as much as shams at zero delay is that their tendency to avoid the delayed reinforcer generalized from trial blocks on which delays were present to the first trial block. Indeed, Figure 92 and Figure 93 show this phenomenon developing.

The task used in the present experiments does not allow the two explanations of impulsive choice (variations in sensitivity to reinforcer magnitude or delay) to be distinguished conclusively. While this may be possible in delay-of-reinforcement choice tasks using indifference-point methodology (Ho *et al.*, 1999, but see Chapter 5), there may be a simpler alternative. Relative preference for two reinforcers is often inferred from the distribution of responses on concurrent VI schedules of reinforcement (see Chapter 1, p. 54). While such an approach is complex when delayed reinforcement is used (see Chapter 1), it is simpler to interpret with immediate reinforcement. If core-lesioned rats were trained on two concurrent VI schedules with identical parameters, with one schedule producing a 1-pellet reward and the other producing a 4-pellet reward, relative preference between the two could be assessed. The matching law (Herrnstein, 1961; 1970) predicts that a subject for whom 4 pellets are worth 4 times as much as 1 pellet would allocate 80% of its responses to the 4-pellet schedule. Normal rats would be expected to perform close to this level, even if they did not 'match' exactly. If core-lesioned subjects exhibited *relative indifference* compared to shams, this would provide independent evidence for reduced reinforcer magnitude discrimination following AcbC lesions (or an abnormality of the matching process itself). If they performed normally, this explanation would become far less likely, in which case the impulsive choice observed in the present experiment could be attributed more specifically to a steeper delay-of-reinforcement gradient.

Published data and the present thesis do not allow this question to be answered directly. However, in Chapter 4, core-lesioned rats were trained on a concurrent VI schedule, albeit for two different reinforcers intended to be of similar value. If anything, these subjects exhibited more pronounced relative preferences than shams (p. 138), indirectly supporting the view that impulsive choice in core-lesioned rats is due to a delay-dependent deficit. However, this issue will require further investigation.

Finally, an explanation in terms of temporal perception might also be offered for the effects of AcbC lesions. The basal ganglia have been suggested to be a component of an 'internal clock', based on the effects of dopaminergic manipulations on timing tasks (see Gibbon *et al.*, 1997). Similarly, forebrain serotonin depletion that affects Acb, among many other structures, impairs timing ability (Morrissey *et al.*, 1993; Wogar *et al.*, 1993a; Morrissey *et al.*, 1994; Al-Zahrani *et al.*, 1997), though these impairments sometimes reflect enhanced behavioural switching rather than a true timing deficit (Ho *et al.*, 1995; Al-Zahrani *et al.*, 1996; Al-Ruwaitea *et al.*, 1997a); see Al-Ruwaitea *et al.* (1997b) for a review. A lesion that increased the speed of an 'internal clock' might (following the distinctions of Killeen & Fetterman, 1988) affect choice prospectively (i.e. the lesioned subject perceives itself to be at a later time-point in the session than it actually is, hastening the within-session shift towards the Immediate lever), or might affect retrospective choice (i.e. the lesioned subject experiences a given delay as longer than it remembered, causing a decrease in its preference for the Delayed lever). Unfortunately, there is at present no evidence to address the question of whether excitotoxic AcbC lesions affect behavioural timing.

Hyperactivity and impulsivity: behavioural comparison to models of ADHD

AcbC-lesioned animals exhibited at least two signs of ADHD: locomotor hyperactivity and impulsive choice (Sagvolden & Sergeant, 1998). However, attentional deficits are not evident in such animals: neither 6-OHDA-induced dopamine depletion of the Acb (Cole & Robbins, 1989) nor excitotoxic lesions of the AcbC (A. Christakou, unpublished observations) affect accuracy in the 5CSRTT test of attentional function.

As discussed above, one possible explanation for the impulsive choices of the AcbC-lesioned group is that these rats were hyposensitive to delayed reinforcement (hypersensitive to the effects of delays). This hypothesis may make predictions about performance on free-operant schedules, discussed below, but first it should be noted that reduced preference for a delayed reward as a *goal* of behaviour in choice experiments is not necessarily the same as reduced ability of delayed reinforcement to strengthen behaviour by 'stamping in' a stimulus–response habit (Thorndike, 1911; Grindley, 1932; Guthrie, 1935; Hull, 1943); goal-directed actions and stimulus–response habits are dissociable (Dickinson, 1994).

Sagvolden *et al.* (1998) suggested that reduced efficacy of delayed reinforcement should lead to hyperactivity (increased responding) on free-operant schedules. For example, subjects responding on FI schedules exhibit a typical 'scallop', in which responding increases as the reinforcer is approached in time; this may be because responses at the end of the interval incur a shorter delay to reinforcement (see Mackintosh, 1974, pp. 170–177). According to this logic, subjects who exhibit a steeper delay-of-reinforcement gradient should show a more pronounced FI scallop, as has been observed for the SHR rat (Sagvolden *et al.*, 1992). However, there are alternative explanations of FI performance (Mackintosh, 1974, pp. 170–171) — indeed, the smooth scallop is only observed when many intervals are averaged, and is not typical of an individual interval (Gentry *et al.*, 1983). Sagvolden *et al.* (1998, p. 62) appear to suggest that the more pronounced scallop is partly a consequence of differential reinforcement of short IRTs; however, it is not clear that this is the case. Ratio schedules do not reinforce particular IRTs with different probabilities, but do reward high rates of responding (short IRTs) with higher local rates of reinforcement, while interval schedules preferentially reinforce *long* IRTs (the longer a subject waits to make the next response, the more likely it is to be reinforced) (see Mackintosh, 1974, p. 177; Dawson & Dickinson, 1990; Tarpy, 1997, pp. 257–258). Regardless of the subject's delay-of-reinforcement gradient, if IRTs represent a basic unit of behaviour to be reinforced (as suggested by Shimp, 1967; 1969), then interval schedules reinforce long IRTs. For the FI scallop to be a consequence of reinforcement of short IRTs, short IRTs would have to occur closer in time to the reinforcer than long IRTs — the scallop phenomenon

intended to be explained. On the other hand, the development of a more pronounced FI scallop is explicable in terms of a steeper delay-of-reinforcement gradient if responses are considered individually.

This issue is of some importance, as it determines whether hyperactivity should follow directly from a steep delay-of-reinforcement gradient. Reduced efficacy of delayed reinforcement does not necessarily imply increased efficacy of immediate reinforcement — Figure 99 illustrates this (compare Figure 1 of Sagvolden *et al.*, 1998). If a steep delay-of-reinforcement gradient does preferentially reinforce short IRTs, it is clear how hyperactivity might emerge (Sagvolden *et al.*, 1998; Sagvolden & Sergeant, 1998). However, if responses are considered individually, the average response would be *less* likely to be reinforced, leading to hypoactivity. Finally, it might be argued that activity levels determine reinforcement efficacy, rather than the other way around. In models such as that of Killeen & Fetterman (1988, p. 288), reinforcement only acts on the behaviour the subject is currently engaged in; delaying the reinforcer simply reduces the probability that the animal has remained in the state associated with that behaviour. In this form of model, changing the rate at which the animal shifts between behaviours — a plausible description of hyperactivity — would be expected to steepen the apparent delay-of-reinforcement gradient.

Thus, there is no clear theoretical compulsion to think that a steep delay-of-reinforcement gradient should produce either hypoactivity or hyperactivity on free-operant schedules. The experimental evidence concerning rats with excitotoxic AcbC lesions indicates that although they exhibit choice behaviour compatible with a steep delay-of-reinforcement gradient (present experiments), and locomotor hyperactivity (present experiments; Maldonado-Irizarry & Kelley, 1995; Parkinson, 1998; Parkinson *et al.*, 1999b), they respond at normal rates on free-operant schedules (e.g. Chapter 4, concurrent VI schedules; Parkinson *et al.*, 1999b, random ratio schedules with conditioned reinforcement).

Implications for theories of nucleus accumbens function

At the least, the present experiments show that the Acb contributes significantly to animals' ability to choose a delayed reward. If further experiments show that it does so specifically by maintaining the value of a reinforcer over a delay, a new avenue of inquiry into Acb function might open up. It has previously been shown in primates that neuronal activity related to the expectation of reward across a delay can be found in the ventral striatum (Schultz *et al.*, 1992; Schultz *et al.*, 1995a; Schultz *et al.*, 1998; Schultz *et al.*, 2000); such activity is a candidate representation of the goals of activity (Schultz *et al.*, 2000). Additionally, striatal neurons may respond to past events, maintaining a form of memory that might assist the association of past acts with reinforcement (Schultz *et al.*, 1995a). These findings represent important data on the forms of information that the AcbC may use to promote actions leading to delayed rewards, and a future challenge will be discover the manner in which these neural signals influence overt behaviour, and the psychological processes they govern. Given the involvement of the Acb in aversive motivation (see Salamone, 1994; Parkinson *et al.*, 1999c), it would also be of great interest to determine whether lesions of Acb induce impulsive choice in an aversive context, impairing the ability to choose a small immediate penalty in preference to a large delayed penalty.

Although the manner in which delayed reinforcement affects free-operant behaviour may be extremely complex, as discussed above, the finding that AcbC lesions reduce subjects' preference for delayed rewards may be useful in interpreting the results of some studies that are at present not clearly understood. For example, Salamone and colleagues have found that dopamine depletion of the Acb leads rats to forgo the opportunity to work for a preferred food, instead consuming more of a less-preferred but freely available food (Salamone *et al.*, 1991; Cousins *et al.*, 1993; Salamone *et al.*, 1994; Cousins *et al.*, 1996), even though reinforcer magnitude discrimination is not overtly impaired by these lesions (Salamone *et al.*, 1994). Similarly, Acb dopamine depletion impairs responding on high-rate but not on low-rate schedules

(McCullough *et al.*, 1993; Salamone *et al.*, 1993; Sokolowski & Salamone, 1998; Aberman & Salamone, 1999). These results have been interpreted as indicating that Acb dopamine depletion impairs the ability of animals to overcome response costs (Salamone, 1994). Although excitotoxic AcbC lesions are clearly not the same as whole-Acb dopamine depletion, two alternative views of these studies are possible. Firstly, as suggested by Parkinson *et al.* (2000a), the impairments may have been due to the loss of a Pavlovian motivational process that normally contributes to instrumental responding (see Chapter 1, p. 50). An interpretation in terms of response costs is certainly inadequate to describe all the data; for example, Acb DA depletion has previously been shown to disrupt displacement behaviours that cannot easily be described as carrying a response cost (Robbins & Koob, 1980). The present results, based on excitotoxic AcbC lesions, provide an even clearer demonstration of the role of the Acb in choice behaviour and the selection of actions, even when those actions do not differ in response effort or cost (in support of Reading *et al.*, 1991; Parkinson *et al.*, 2000a). A second interesting interpretation of the results of Salamone and colleagues, based on the present data, is that the lesions reduced the subjects' inclination to respond for food, particularly on high-rate schedules, because that reward was significantly *delayed*. Instead, the lesioned rats preferred an immediately-available but smaller reward.

Effects of intra-accumbens amphetamine

Theories that attribute impulsive choice to hypofunctional Acb DA systems (see Sagvolden & Sergeant, 1998) thereby suggest that Acb DA normally contributes to the effectiveness of delayed reinforcement (and thereby to self-controlled choice), and would predict that intra-Acb amphetamine would increase preference for delayed reward. Yet the opposite was observed. Injections of amphetamine into the Acb reduced subjects' preference for the large, delayed reward slightly, but not in a clear dose-dependent manner; over all subjects, the 3- μ g and 20- μ g doses had this effect, but the 10- μ g dose did not differ from saline. Furthermore, despite the prediction made in Chapter 6 (p. 192) that intra-Acb amphetamine might enhance the effects of cue stimuli present during a delay to reinforcement to promote 'self-controlled' choice, no cue-dependent effects of amphetamine were observed, and the effects of amphetamine were reasonably consistent across the two groups (Figure 97, p. 222).

A new version of the task was used for this experiment. The sessions were shorter and the delays were arranged in reverse order (with the longest delay presented at the start of the session), in an attempt to ensure that high Acb levels of drug coincided with responding at non-zero delays. However, this new task may have produced methodological problems. The task appeared more difficult for subjects to acquire than the version used in other experiments, with lower levels of preference for the large reinforcer at zero delay (compare acquisition in Figure 96, p. 221, with that in Experiments 1/2/3 and Chapter 6) and more pronounced differences in absolute preference levels between the Cue and No Cue groups during acquisition (though not in the slope of the within-session shift in preference). If subjects are to reach the same levels of preference at each delay as in the 'standard' version of the task, their within-session shift in preference must be more rapid; possibly this is harder to learn. Both this and the more pronounced group differences in responding may have rendered the abbreviated task less sensitive to pharmacological manipulations.

One other methodological issue is the order in which the delays were given. When drug effects are tested with only an ascending, or only a descending, series of delays, any delay-dependent effects of the drug are confounded with the pattern of responding across a session. It was therefore hoped that the use of a descending series of delays would allow some comparison with the effects of systemic amphetamine observed in Chapter 6, with the potential to distinguish (for example) a tendency to complete the within-

session shift in preference more rapidly from a true effect on preference for delayed reward. However, this training technique meant that the first time the subjects experienced a choice of the two levers, one lever delivered a larger reinforcer but after a 60-s delay. For a subject accustomed to continuous reinforcement, this may have induced rapid extinction on that lever, an effect that might have contributed to poor learning in the present experiment.

In the absence of a clear dose-dependent effect of amphetamine on choice, or on other measures of performance, it is difficult to draw firm conclusions. Taken at face value, the present results indicate that intra-Acb amphetamine causes a slight reduction in preference for delayed rewards, without affecting motoric aspects of task performance, and that signals present during the delay do not contribute to its action at this site. The Acb might therefore be a neural locus of the cue-independent effects of systemic amphetamine (see Chapter 6), but the locus of the cue-dependent effects remains uncertain. However, not only was it possible that the new task was relatively insensitive to the effects of amphetamine (discussed above), but the injector tips in this experiment were on the whole located more ventrally than was intended (Figure 95, p. 219); amphetamine was injected into ventral AcbSh and/or underlying structures, and this fact may also account for the lack of a systematic effect of amphetamine on choice. It will probably prove worthwhile to replicate this experiment comparing the effects of amphetamine injections in the ventromedial AcbSh with injections in the AcbC, particularly given the newly-discovered role of the AcbC in preference for delayed reward (Experiment 3). In doing so, it may also help to adjust the task parameters in an attempt to avoid some of the pitfalls discussed here (ensuring that the abbreviated task is sensitive to the behavioural and systemic pharmacological manipulations used in Chapter 6), or simply to use intra-Acb amphetamine with the full, 100-min task.

Finally, it is interesting to note that rats reared in isolation have recently been found to be less impulsive than socially-reared controls on the delayed reinforcement choice task used in the present experiments (full version, without signals during the delay), and this difference was exaggerated by systemic *d*-amphetamine (Y.-P. Liu, L.S. Wilkinson and T.W. Robbins, unpublished observations; L.S. Wilkinson, personal communication, 4 January 2001). Isolation-reared rats exhibit augmented Acb DA release in response to psychostimulant drugs (Jones *et al.*, 1992; Howes *et al.*, 2000), with some studies showing elevated basal levels of Acb DA (Hall *et al.*, 1998), but they also exhibit other neurochemical abnormalities, including differences in 5-HT levels in the Acb and DA levels in the mPFC and amygdala (Jones *et al.*, 1992; Heidbreder *et al.*, 2000). These differences represent other candidate systems where anatomically- and neurochemically-specific drug infusions might affect impulsive choice.

Autoshaping and impulsivity

Autoshaping itself has been suggested to reflect impulsive behaviour, in that subjects are unable to withhold responses to the CS (Tomie, 1996). Subjects' propensity to autoshape has previously been shown to predict sensitivity to delays in a similar delay-of-reinforcement procedure to that used here (Tomie *et al.*, 1998). Individuals that autoshape readily have been suggested to be more vulnerable to drugs of abuse (Tomie, 1996), while impulsive choice behaviour predicts alcohol self-administration in rats (Poulos *et al.*, 1995). Rats that autoshape readily have higher levels of dopamine and dopamine metabolites in the Acb than rats that do not (Tomie *et al.*, 2000), while dopamine depletion of the Acb and excitotoxic lesions of the AcbC both impair autoshaping (Everitt *et al.*, 1999; Everitt *et al.*, 2000b; Parkinson *et al.*, 2000c; Parkinson *et al.*, submitted).

However, the relationship between impulsivity and autoshaping has not been clearly established. Autoshaping is suggested to represent impulsivity in that the subject is unable to suppress the involuntary

tendency to approach the CS — ‘motor impulsivity’, or failure of inhibitory control (Tomie, 1996; Tomie *et al.*, 1998). Motor impulsivity has been doubly dissociated from impulsive choice by pharmacological means (summarized by Evenden, 1999b). The correlation between the two suggested by Tomie *et al.* (1998) is therefore not a trivial result. Unfortunately, in the study of Tomie *et al.* (1998), which used a task very similar to that of Evenden & Ryan (1996), over 50% of the subjects showed exclusive preference, choosing the large delayed reward or the small immediate reward at all delays. These subjects scored zero on the measure of impulsive choice used by Tomie *et al.* for correlation with autoshaping CR frequency. Furthermore, the autoshaping stimulus was almost identical to one of the levers used subsequently in the delayed reinforcement choice task; thus, autoshaping and impulsive choice may have been correlated not because of an underlying common cause (impulsivity), but because differences in subjects’ experience with the autoshaping stimulus affected choice directly. Finally, the autoshaping task used included no control stimulus (CS–) unpaired with reward; therefore, autoshaping performance in their study was potentially confounded with differences in unconditioned behaviour. (Different autoshaping tasks may also generate different views of what constitutes motor impulsivity: if a subject responds to the CS+ but not to a similar CS–, is it showing good impulse control by suppressing responses to the CS–, or poor impulse control by responding to the CS+ in the first place? One view is that total CS responding is an index of impulsivity, in which case an absence of responding indicates good impulse control, selective CS+ responding indicates mild impulsivity, and responding to both the CS+ and the CS– indicates grossly impaired impulse control.) While Tomie *et al.*’s (1998) result appears to indicate that *sensitivity* to the delays in the choice task correlate with either the propensity to autoshape, general activity, or exploratory tendencies, it is not clear that simple sensitivity to delays is the same as impulsive choice. In particular, it is not obvious that subjects who always chose the small immediate reward in this task exhibited ‘zero impulsivity’, and the proportion of rats exhibiting this insensitivity to delay may reflect procedural differences (such as the method of training, as suggested in Chapter 6, p. 190). Future investigations of this important area will need to pay particular attention to the definitions of impulsivity used.

The present study raises two further dissociations between autoshaping and impulsive choice. First, lesions of the ACC are known to impair autoshaping, generally in the ‘disinhibited’ fashion of increasing approaches to a neutral CS– (Chapter 3; Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c). Of course, this may represent a different idea of the relationship between motor impulsivity and autoshaping than that of Tomie *et al.* (1998); ACC lesions have also been suggested to increase motor impulsivity via disinhibition in other tasks (Muir *et al.*, 1996) (though these lesions may differ slightly; see p. 119 and Figure 14, p. 72). However, ACC lesions did not affect impulsive choice in the present experiments. Second, lesions of the AcbC, which abolish the development and performance of autoshaping by reducing approaches to the CS+ (Everitt *et al.*, 2000b; Parkinson *et al.*, 2000c), rendered rats dramatically *more* likely to make impulsive choices.

The possible role of other structures connected to the nucleus accumbens core

It has been shown that while lesions of the AcbC impair rats’ capacity to choose a delayed reward, lesions of two of its afferents did not (mPFC lesions produced a deficit but this was qualitatively different). An important task for further investigations is to specify which afferents to the AcbC contribute to its ability to promote the choice of delayed rewards, and through what efferent pathways it does this.

One obvious afferent structure that may provide specific information concerning reinforcer value to the Acb is the BLA, while the CeA might affect preference by modulating the dopamine innervation of the Acb. Another direct afferent is the orbitofrontal cortex, also implicated in the assessment of reward

value and probability (Rogers *et al.*, 1999) (see also Chapter 1 for a discussion of amygdala and orbitofrontal cortex function, and Öngür & Price, 2000 for the delineation of the orbitofrontal cortex in the rat). The orbitofrontal cortex may also be an important efferent target of information travelling through Acb, as this 'limbic loop' of the basal ganglia projects back (through the ventral pallidum) to medial orbitofrontal cortex (Alexander *et al.*, 1986). In addition, it remains to be seen whether the AcbSh also plays a role in the choice of delayed rewards. This is another interesting target of investigation, given the abnormalities of dopamine receptor function detected in the AcbSh of the SHR (Papa *et al.*, 1996; Carey *et al.*, 1998; Papa *et al.*, 1998; Sadile, 2000).

Finally, the limbic corticostriatal circuit may not be the only system involved in delayed reinforcement. In principle, any structure that represents *future reinforcers* across a delay may contribute to the choice of future reinforcers, and exert conditioned reinforcing effects on current behaviour, while any structure that maintains a 'memory trace' of responses across a delay may support the reinforcement of *past responses*. The ventral striatum and orbitofrontal cortex exhibit such activity (Schultz *et al.*, 1995a; Schultz *et al.*, 1998; Schultz *et al.*, 2000), but so do other structures including the dorsal striatum (e.g. Schultz *et al.*, 1995a), implicated in the reinforcement of stimulus–response habits (see Chapter 1, p. 46).

Conclusions

The present results provide direct evidence to support previous conjectures that the Acb is involved in the pathogenesis of impulsive choice. Hitherto, these conjectures have been based on correlational data, including findings of neurochemical abnormalities in the Acb of animal models of ADHD (see Sagvolden & Sergeant, 1998); the present study demonstrates a causal role for Acb dysfunction in impulsive choice. No evidence was found for similar involvement of the ACC or mPFC. It remains to be seen whether failure of Acb dopamine function can also contribute to impulsive choice. The remainder of the neural circuit underlying the efficacy of delayed reinforcers remains to be elucidated, but the present methodology holds promise as a means of identifying it.

Chapter 8.

General discussion

Introduction

The experiments described in this thesis addressed the role played by regions of the prefrontal cortex and ventral striatum in the control of rats' behaviour by Pavlovian conditioned stimuli, and in their capacity to choose delayed reinforcement. In this concluding chapter, the findings from these experiments will first be summarized briefly. The results have already been discussed in Chapters 3–7; in this chapter, their implications will be considered in a wider context and future research directions will be suggested. The role of the ACC within its corticocortical and corticostriatal circuits will be discussed first in the light of the present data. Different theoretical views of the process of choosing between delayed rewards will then be considered, together with the neural basis of this process. Implications for theories of nucleus accumbens function will be discussed, and lastly an overview of reinforcement processes will be presented.

Summary of results

Role of the anterior cingulate cortex in Pavlovian conditioning

The ACC has previously been strongly implicated in stimulus–reinforcer learning in the rodent, in both appetitive (Bussey *et al.*, 1996; 1997a; 1997b; Parkinson *et al.*, 2000c) and aversive settings (Gabriel *et al.*, 1980a; Gabriel *et al.*, 1980b; Buchanan & Powell, 1982a; Gabriel & Orona, 1982; Gabriel *et al.*, 1991a; Gabriel *et al.*, 1991b; Gabriel, 1993; Powell *et al.*, 1994). In Chapter 3, rats with excitotoxic ACC lesions were tested on a variety of tasks to which stimulus–reinforcer learning was expected to contribute. Lesioned rats were impaired at the acquisition of autoshaping, replicating previous findings (Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c), and were also impaired when the lesion was made following training. Unexpectedly, however, they were unimpaired on a number of other tasks based on Pavlovian conditioning procedures and encompassing a range of behavioural responses. ACC-lesioned rats performed normally on a simple temporally discriminated approach task, and responded normally for a conditioned reinforcer (with normal potentiation of this responding by intra-accumbens amphetamine). They also exhibited normal conditioned freezing to an aversive CS, and normal PIT. However, ACC-lesioned rats were impaired on a two-stimulus discriminated approach task (designed to capture features both of autoshaping and the conditioned approach task on which they were unimpaired), providing direct support for the hypothesis that the ACC is critical for discriminating multiple stimuli on the basis of their association with reward.

Role of the nucleus accumbens core and shell in response-specific Pavlovian-instrumental transfer

It has previously been shown that the AcbC contributes to simple PIT (Hall *et al.*, 1999). In Chapter 4, the contribution of the AcbC and AcbSh to response-specific PIT was assessed; this more complex task involves the direction of instrumental choice behaviour by noncontingently-presented Pavlovian CSs. Lesions of the AcbC impaired the response specificity of PIT (that is, the ability of the CS to influence choice behaviour) while lesions of the AcbSh impaired PIT itself. These results present problems of interpretation in the light of other studies, discussed in Chapter 4, but closely resemble the effects of AcbC

and AcbSh lesions on the effects of intra-Acb psychostimulants on responding for conditioned reinforcement (Parkinson *et al.*, 1999b), with the shell providing ‘vigour’ and the core ‘direction’ for PIT.

Behavioural tasks used to assess preference for delayed reinforcement

In Chapters 5 & 6, two tasks testing subjects’ ability to choose a large, delayed reward in preference to a small but immediate reward were investigated in detail. In Chapter 5, rats were tested on a version of the adjusting-delay schedule (1987; Mazur, 1988; 1992; Wogar *et al.*, 1993b). Surprisingly, no direct evidence was found that the subjects were sensitive to the contingencies operating in this schedule, despite the use of a novel cross-correlational analysis that successfully detected such sensitivity in a range of computer-simulated subjects. For this and other reasons, this task was not pursued further. Instead, in Chapter 6, a modified version of the ‘systematic’ technique of Evenden & Ryan (1996) was considered. Using this task, it was demonstrated that rats were directly sensitive to the delay to reward, preferring a large reward less when it was delayed. In a detailed behavioural analysis of the task, the effects of extinction, delay omission, reversal of the pattern of delays presented to the subjects, and satiation were examined, together with the effects of cues present during the delay to reward, thereby partially characterizing the basis of normal subjects’ performance. In particular, it was found that if subjects were trained with a signal or cue present during the delay to the large reward, the cue speeded learning and supported choice of the large reinforcer.

Effects of d-amphetamine, α -flupenthixol, and chlordiazepoxide on preference for signalled and un-signalled delayed reinforcement

In Chapter 6, groups of rats were trained on the delay-of-reinforcement choice task, with or without an explicit signal present during the delay. *d*-Amphetamine, α -flupenthixol, and chlordiazepoxide were then administered before their performance was again tested. Amphetamine enhanced preference for the large, delayed reward in the presence of the cue, at certain doses, but uniformly depressed this preference in subjects trained without the cue. This was suggested to reflect the known effect of amphetamine to enhance the efficacy of conditioned reinforcement (Hill, 1970; Robbins, 1976; Robbins, 1978; Robbins *et al.*, 1983), and may explain discrepancies in the literature regarding the effects of amphetamine on impulsive choice (Evenden & Ryan, 1996; Richards *et al.*, 1997a; 1999; Wade *et al.*, 2000). Flupenthixol, known to depress responding for conditioned reinforcement (Robbins *et al.*, 1983; Killcross *et al.*, 1997a), had cue-dependent effects consistent with this hypothesis, though it generally decreased preference for the delayed reward. The effects of chlordiazepoxide, a benzodiazepine expected not to affect conditioned reinforcement (Killcross *et al.*, 1997a), did not depend on the cue condition: chlordiazepoxide generally reduced preference for the delayed reward.

Neural basis of preference for delayed reinforcement

In Chapter 7, the same delayed-reinforcement choice task was used to assess the contribution of subregions of the ventral striatum and prefrontal cortex to preference for delayed reward. Subjects were trained on the task in the absence of explicit cues, matched to groups, and received sham surgery or lesions of the ACC, mPFC, or AcbC before being retested. ACC lesions had no effect on choice behaviour, though lesioned subjects were slower to collect the large, delayed reward. Lesions of the mPFC altered choice, but not in a manner interpretable as an altered effect of the delays. Rather, mPFC-lesioned rats exhibited a ‘flattening’ of the within-session shift from the large to the small reward as the large reward was progressively delayed; this was suggested to reflect a loss of session-wide temporal stimulus control. In contrast, lesions of the AcbC dramatically and persistently impaired subjects’ ability to choose the delayed reward,

even though the subjects discriminated the two reinforcers. In a new, abbreviated version of the task, infusion of amphetamine into the Acb reduced subjects' preference for the delayed reward, but surprisingly did not do so in a clear dose-, delay-, or cue-dependent manner.

The results of the lesion studies reported in this thesis may be integrated into other work within this field as shown in Table 21 (overleaf).

Anterior cingulate cortex function

The relationship of the present findings to other theories of rodent and primate ACC function were discussed in Chapter 3, in which it was suggested that the rat ACC 'disambiguates' similar stimuli for its corticostriatal circuit on the basis of their differential association with reinforcement. It has been shown that the ACC–AcbC projection is necessary for rats to acquire the autoshaping task used in the present experiments (Parkinson *et al.*, 2000c). As lesions of the AcbC also impair conditioned approach in a temporally discriminated approach task (Parkinson *et al.*, 1999b), suggesting a general role for AcbC in conditioned approach, it would be predicted that ACC–AcbC disconnection would impair the acquisition of the two-stimulus temporally discriminated approach task developed in Chapter 3. This hypothesis awaits experimental test.

The ACC provides specific information to the Acb via glutamatergic projections, through which it influences response selection in conditioned approach tasks (Parkinson *et al.*, 2000c), just as the BLA appears to do for conditioned reinforcement (Burns *et al.*, 1993) and probably for PIT (Blundell & Killcross, 2000a). In all these tasks, the glutamatergic information is in some manner 'gated' or amplified by the dopaminergic innervation of the Acb, probably under the control of the CeA (Cador *et al.*, 1991; Robledo *et al.*, 1996; Hall *et al.*, 1999; Parkinson *et al.*, 2000b; Parkinson *et al.*, submitted). On the basis of other studies reviewed in Chapters 1 & 3, it is suggested that the contributions of the BLA and ACC differ in the following way: the BLA uses a CS to retrieve the motivational value of its specific US, while the ACC directs responding on the basis of the specific CS, preventing generalization to similar CSs. These suggested roles are different — the contributions of the two structures have been dissociated using autoshaping (Bussey *et al.*, 1997a; Parkinson *et al.*, 1999a) and conditioned reinforcement tasks (Chapter 3; Burns *et al.*, 1993) — but are not dissimilar, and it may be a promising area for future research to determine how these two interconnected structures communicate, and the function of that communication.

Additionally, the results of Chapter 7 provide evidence that the ACC is not simply required when behavioural tasks become 'difficult'. In the delayed reinforcement choice task, the within-session increase in the delay to the large reward causes a progressive decline in normal subjects' success at obtaining food. This can plausibly be interpreted as an increase in task difficulty, yet ACC lesions did not impair performance. As a general role for the ACC in 'task difficulty' is an untenable interpretation, further support is inferred for the specific hypothesis that *the ACC is a reinforcement learning structure involved in stimulus discrimination*. The results of Chapter 3 are also not parsimoniously explained by a deficit in *response* discrimination, as the two-stimulus temporally discriminated approach task measured exactly the same response following presentation of a CS+ or a CS–; thus, no response discrimination was required, and yet a deficit was still observed in ACC-lesioned animals.

Table 21. Summary of lesion studies concerning the major tasks used in this thesis, and related work. Results from this thesis are emboldened; a dash (–) indicates no data are available. References (* indicates studies that have not been peer-reviewed fully): (1*) Cardinal, this thesis; (2) Bussey *et al.* (1997a); (3) Parkinson *et al.* (2000c); (4) Burns *et al.* (1993); (5) Parkinson *et al.* (2000b); (6*) Everitt *et al.* (2000b); (7*) Parkinson *et al.*, submitted; (8) Cador *et al.* (1989); (9) Parkinson *et al.* (1999b); (10*) Hall *et al.* (1999); (11) Killcross *et al.* (1997b); (12) Robledo *et al.* (1996); (13*) Killcross *et al.* (1998); (14*) Coutureau *et al.* (2000); (15*) Blundell *et al.* (2000a); (16*) Dix *et al.* (2000); (17) Taylor & Robbins (1986); (18*) Corbit & Balleine (2000a); (19) Morgan & LeDoux (1995); (20) reviewed by e.g. LeDoux (2000); (21) but see Parkinson *et al.* (1999c).

<i>Effects of excitotoxic lesions to:</i>	ACC	mPFC	BLA	CeA	AcbC	AcbSh	Acb DA depletion
Approach tasks							
autoshaping (acquisition)	impaired ^{1,2,3}	normal ²	normal ⁵	impaired ⁵	impaired ³	normal ³	impaired (severely) ⁷
autoshaping (performance)	impaired ¹	–	–	normal ⁶	impaired ¹	–	impaired (mildly) ⁷
temporally discriminated approach (Pavlovian)	normal ¹	–	normal ⁸	normal ¹²	impaired ⁹	normal ⁹	–
discriminated approach (instrumental contingency)	–	normal ⁴	impaired ⁴	–	–	–	–
discriminated approach (Pavlovian, two stimuli)	impaired ¹	–	–	–	–	–	–
General potentiation/suppression of instrumental behaviour by a conditioned cue							
simple Pavlovian–instrumental transfer (conditioned elevation)	normal ¹	–	normal ¹⁰	impaired ^{10,13}	impaired ¹⁰	normal ¹⁰	–
conditioned suppression	–	normal ¹⁴	normal ¹¹	impaired ¹¹	normal (whole Acb lesion) ¹⁶	–	–
intra-Acb amphetamine potentiation of CRF	normal ¹	normal ⁴	present/altere ⁴	impaired ¹²	loss of specificity ⁹	impaired ⁹	impaired ¹⁷
Directed modulation of instrumental behaviour by a conditioned cue							
conditioned punishment	–	impaired ¹⁴	impaired ¹¹	normal ¹¹	impaired (whole Acb lesion) ¹⁶	–	–
conditioned reinforcement	normal ¹	normal ⁴	impaired ^{4,8,13}	normal ^{8,13}	normal ⁹	normal ⁹	normal ¹⁷
response-specific Pavlovian–instrumental transfer	–	–	impaired (loss of specificity) ¹⁵	–	impaired (loss of specificity)¹; normal¹⁸	impaired (loss of transfer)^{1,18}	–
Other Pavlovian conditioning procedures							
conditioned freezing to a discrete CS	normal ¹	normal/enhanced ¹⁹	impaired ²⁰	impaired ²⁰	– ²¹	– ²¹	– ²¹
Delayed reinforcement							
Ability to choose a large, delayed reinforcer over a small, immediate reinforcer	normal ¹	intact, though loss of usual pattern of responding¹	–	–	impaired ¹	–	–

There have been several suggestions that ACC dysfunction is related to impulsive behaviour or over-responding (Muir *et al.*, 1996; Bussey *et al.*, 1997a; Parkinson *et al.*, 2000c). However, Chapter 7 demonstrated that ACC lesions do not induce impulsive choice, in addition to providing evidence for a behavioural dissociation between autoshaping and impulsive choice through lesion studies of the ACC and AcbC. Over-responding to a CS– in a task such as autoshaping may reflect a failure of discrimination, rather than impulsive responding. Thus, to investigate whether the ACC is truly involved in impulsivity in any way, explicit tests of motor impulsivity (such as a paced fixed consecutive number schedule, in which subjects must avoid terminating chains of responses prematurely, or a ‘stop’ task, in which subjects must inhibit ongoing behaviour) or reflection impulsivity (failure to acquire sufficient information to perform a task accurately) (see Evenden, 1999b) should be administered to subjects with ACC lesions.

Finally, one of the most interesting questions about the function of the ACC concerns its apparently time-limited role in behaviour (see Chapter 3, p. 113). This makes analysis of its function more difficult, as it is not presently possible to predict accurately when in the course of behavioural training the contribution of the ACC is no longer significant. As Chapter 3 also made clear, this issue touches on the present boundaries of understanding of the way in which the representations formed during Pavlovian conditioning change over time. There are several critical issues. (1) Can overtrained Pavlovian responding be considered habitual? (2) With what structures does the ACC interact during learning, and how? Candidates include the Acb, amygdala, OFC, and PCC. The ACC may do more than simply provide a flexible behavioural controller that is effective while other structures are learning more permanent representations, but it may also actively ‘teach’ other structures such as the PCC (Gabriel *et al.*, 1980a, p. 162; Gabriel, 1993; Freeman *et al.*, 1996; Hart *et al.*, 1997). This hypothesis provides a testable prediction concerning appetitive autoshaping: that well-learned performance will be sensitive to PCC lesions (see Chapter 3, pp. 99/113) even though early acquisition is not (Bussey *et al.*, 1997a).

Theories of learning and choice with delayed reward

Two broad approaches to choice behaviour will be summarized, and a synthesis offered.

Model 1 (informed choice). According to this model, subjects make prospective choices between alternatives based on full knowledge of the response–outcome contingencies and of the value of each outcome. These choices represent goal-directed actions. Subjects’ sensitivity to delay in choice tasks is therefore a consequence of time discounting of the perceived (prospective) value of the delayed reward.

This model is necessarily applicable only to fully-trained subjects — subjects who have learned the instrumental contingencies. It may be particularly applicable when humans are offered explicit hypothetical choices (‘would you prefer \$800 now, or \$1000 in a year?’; Rachlin *et al.*, 1991; Myerson & Green, 1995).

As the contingencies cannot be offered ‘pre-packaged’ to experimental animals through language, such subjects must be trained through direct experience of the rewards in the experimental situation. This introduces the complication that delays to reinforcement can affect operant and discrimination learning (reviewed in Chapter 1), so care is typically taken by experimenters to ensure subjects are ‘well trained’. Slow acquisition of delay sensitivity must be attributed to difficulties in learning the instrumental contingencies across a delay and/or learning the appropriate incentive value of delayed reward through experience of waiting. In tasks where the delay is systematically and predictably varied, as in Chapters 6 & 7, learning may also be slowed by the requirement to learn S^D s predicting the delay contingency currently in

force. Thus, this model is inherently an incomplete description of the effects of delayed reinforcement, as it does not deal with the effects of delays on learning.

Model 2 (associative response strength). According to an extreme form of this model, based on simple S–R theory (Thorndike, 1911; Grindley, 1932; Guthrie, 1935; Hull, 1943), rats' choice behaviour reflects differential reinforcement of stimulus–response habits. The change in associative strength is some function of reward magnitude multiplied by the time-discounted 'trace strength' of the preceding response. Choice is determined by some process of competition between the available responses (e.g. the principles of matching; Herrnstein, 1970; de Villiers & Herrnstein, 1976). Choice is therefore 'retrospective' in a sense, as preference for a particular alternative depends upon prior experience of that alternative, and time discounting reflects the decay of the traces available to be associated with reward. A similar model, after Grice (1948), may be constructed in which animals respond for immediate conditioned reinforcement (by goal-directed behaviour or S–R habit) and the acquisition of associations between a chain of stimuli and eventual reward accounts for the observed phenomenon of temporal discounting, by similar mechanisms.

The S–R view accounts for some of the theoretical appeal of exponential temporal discounting models. In exponential decay, at any one moment in time the trace strength of a response follows directly from the trace strength at the previous instant (if x_t is the trace strength at time t and A is the starting value, then $x_t = Ae^{-kt}$ and $x_{t+1} = e^{-k}x_t$). In contrast, in the hyperbolic discounting model and all others in which preference reversal occurs, the strength of the trace at any one moment cannot be calculated in such a manner: information about the absolute time since the response must be available. (This may be clearly illustrated by the preference reversal graph shown in Chapter 1, p. 58; if two such decay curves cross, then an observer travelling along one curve cannot know at the crossover point whether its own curve is the recent, rapidly-decaying trace, or the older, slowly-decaying trace, without further information — namely the time since the response or its starting strength.) This process does not model 'mnemonic delay' in any clear way. Thus, the empirical observation of hyperbolic discounting specifies the information that must be available to the subject at any one moment in time; in the context of 'retrospective' choice, this constrains the possible underlying psychological mechanisms, and there is no obvious candidate within the S–R model.

While S–R models can account for effects of delays on learning as well as choice, they do not take into account the fact that goal-directed actions contribute to choice in rats (Dickinson, 1994) and would clearly not provide a satisfactory account of human choice (cf. Ainslie, 1975; Rachlin *et al.*, 1991; Myerson & Green, 1995).

Model 3 (composite). A multifactorial model is therefore suggested, based on that of Dickinson (1994). The 'response strength' of any behaviour is governed by (1) goal-directed action (Dickinson & Balleine, 1994), in which knowledge of the instrumental contingency combines with the incentive value of the expected outcome; (2) stimulus–response habits, which gain strength slowly with the number of reinforcers presented (Dickinson *et al.*, 1995); and (3) PIT, mediated by the Pavlovian association between contextual, discriminative, or other conditioned stimuli and the outcome of the instrumental action. Ordinarily, behaviour conforming to the matching law and to hyperbolic temporal discounting is seen as a product of these processes. Delayed reinforcement may act (a) to impair learning of the instrumental contingency (Dickinson *et al.*, 1992); (b) to reduce the incentive value of the delayed reward, as speculated by many models; (c) to reduce the reinforcement of stimulus–response habits; and (d) to reduce the Pavlovian association between stimuli present at the moment of action and the ultimate reinforcer.

This model makes several predictions. Firstly, manipulations of components of this composite behaviour should affect choice. For example, manipulations of the association between cues immediately con-

sequent on choice and the outcome (e.g. presence of absence of a cue bridging the delay) should affect choice independently of the actual delay to reinforcement, a prediction not made by Kacelnik's (1997) normative model of hyperbolic discounting, but one supported by the results of Chapter 6. Secondly, pharmacological and neural manipulations known to dissociate these processes should also be capable of affecting choice.

This view is obviously compatible with mathematical models of temporal discounting, but interprets the discount function as the sum of the contributions of several processes operating in any one situation. Similar composite models have been offered before (a casual example is Pinker, 1997, pp. 395–396), though with a different decomposition of the processes contributing to choice (e.g. distinct contributions of conditioned and primary reinforcement to response strength; Killeen & Fetterman, 1988, pp. 287–289). One interesting challenge may be to establish what processes contribute most significantly to choice of a reinforcer at different delays. Consider an obvious hypothesis: instrumental incentive value in the rat depends upon declarative knowledge, as discussed in Chapter 1 (p. 23), and in this way is analogous to human hypothetical choices. Thus it may be that when reward is extremely delayed (as in some human experiments), only instrumental incentive value is important (as delay $d \rightarrow \infty$, total value $V \rightarrow V_{\text{instrumental}}$). When a dieting human calmly decides to abstain from chocolate cake and the dessert trolley is then pushed under his nose, it would not be expected from the rat literature that the instrumental incentive value of chocolate cake suddenly changes — after all, the subject's underlying motivational state of hunger (or lack of it) has not altered. However, alternative, possibly Pavlovian motivational processes may create an extra boost to the value of the cake (observed as a tendency to choose the cake), which is now immediately available (as $d \rightarrow 0$, $V_{\text{cake-other}}$ increases dramatically). The net value function ($V = V_{\text{cake-instrumental}} + V_{\text{cake-other}}$) could then exhibit preference reversal, leading our diner to succumb and choose the immediate reinforcer. This illustrates but one possible scenario. Nevertheless, if different processes do contribute at different delays, there would be important implications for our understanding of individual differences in impulsive choice.

Devaluation of the delayed reinforcer may yield clues concerning this suggestion in experimental animals. For example, if devaluation led to a fall in preference for a delayed reward (relative to a non-devalued, immediate alternative) when subjects were tested in extinction, this would suggest that instrumental incentive processes were prominent contributors to the overall 'value' of the delayed reward, while failure to observe this would suggest habitual responding. If instrumental incentive processes contribute more to the value of an immediate reinforcer than to that of a delayed reinforcer, and the two reinforcers were of the same foodstuff, devaluation might even lead to an increase in preference for the delayed reinforcer (assuming the subjects did not simply cease responding). Of course, a practical problem might be that the use of a discrete-trial schedule may encourage stimulus-bound responding in a way that free-operant schedules do not, while providing frequent choices may discourage habit formation.

Theories of nucleus accumbens function and the neural basis of delayed reward

1. *The striatum as a switching device*

A quarter of a century ago, Lyon and Robbins (1975) hypothesized a behavioural switching mechanism based on the dopaminergic innervation of the striatum. This concept has evolved (Robbins & Sahakian, 1983; Robbins *et al.*, 1990b); Redgrave *et al.* (1999a) recently reviewed and extended theories of the basal ganglia as a central behavioural switching mechanism (Lyon & Robbins, 1975; Cools, 1980; Dunnett & Iversen, 1982; Jaspers *et al.*, 1984; Redgrave *et al.*, 1999a), which provides a useful framework within

which to discuss the present data (see also Parkinson *et al.*, 2000a). According to this theory, the striatum selects responses in the cortical structures to which it is connected, by disinhibiting one ‘channel’ passing through it, and using a winner-take-all system to ensure that only a single channel is active. Superimposed upon this picture may be a hierarchy: whilst the motor loop of the dorsal striatum switches between incompatible commands to the musculature, the limbic loop (ventral striatum) may operate at a higher level to switch between different overall behavioural strategies. The concept of hierarchical switching is illustrated in Figure 100 (with detail of the ventral striatal circuit in Figure 101). This mechanism is an efficient way to resolve conflicts over access to limited motivational, cognitive and motor resources (Redgrave *et al.*, 1999a).

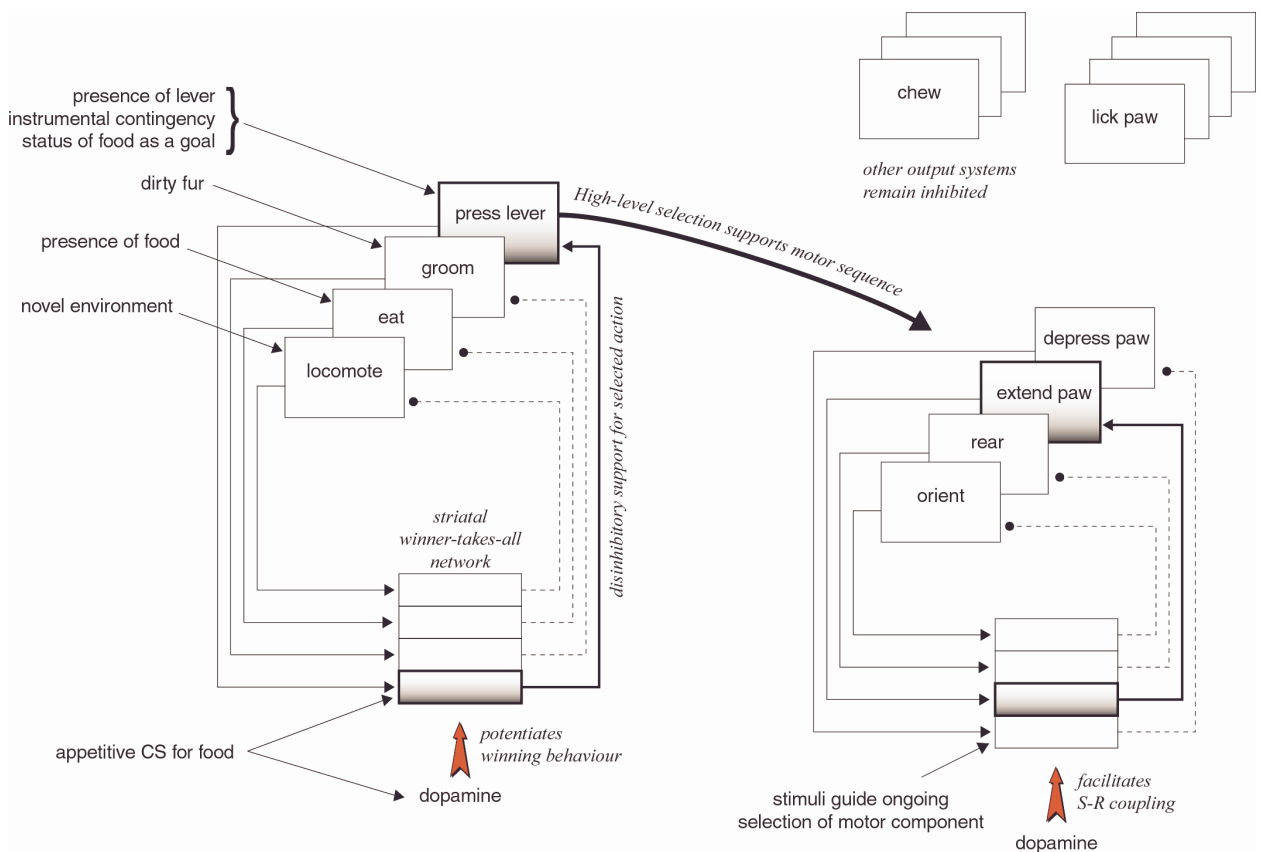


Figure 100 illustrates concepts of central switching mechanisms and hierarchies of behaviour. The left-hand circuit, representing the limbic corticostriatal loop, influences the selection of complex behaviours on the basis of conditioned motivational stimuli. The right-hand circuit, representing the motor corticostriatal loop, selects motor responses on the basis of environmental stimuli in an S–R fashion. The interaction between the circuits represents the hierarchy of behaviour: motor components can only be selected when they are part of the chosen higher-level behaviour.

Striatal circuitry is consistent with this hypothesis. Striatal medium spiny neurons are well suited by their connectivity and electrophysiology to act as pattern detectors: they are bistable, receive a highly convergent projection from the cortex and require cortical input to enter the active (‘up’) state. They are therefore suited to ‘registering’ patterns of cortical input (see Houk & Wise, 1995; Wilson, 1995) and appear to do so (Schultz *et al.*, 1995a). More controversially, they may receive a ‘teaching signal’ to influence future recognition of cortical patterns of activity, discussed later. A caveat is that the neostriatum is only able to discriminate cortical input patterns that are linearly separable, as it is equivalent to a single-layer network (Wickens & Kötter, 1995), and its discriminative ability is further limited by the fact that direct




Figure 101. Speculative view of influences mediated through the ventral striatum ('close-up' of the left-hand circuit of Figure 100). Information about CSs may influence the Acb in several ways. The BLA is implicated in the retrieval of the current value of the US (see Everitt *et al.*, 2000a), and the modulation of choice behaviour; for example, it is required for CRf (Cador *et al.*, 1989; Burns *et al.*, 1993) and response-specific PIT (Blundell & Killcross, 2000a); the latter requires the AcbC (Chapter 4). Among its other roles, the CeA projects to the VTA and is required for the invigorating effect of CSs on instrumental responding (in tasks such as simple PIT; Hall *et al.*, 1999) and on locomotor approach (in autoshaping; Parkinson *et al.*, 2000b), probably via its effects on Acb DA (Hall *et al.*, 1999; Parkinson *et al.*, submitted). The ACC appears to be required to discriminate similar CSs on the basis of their association with reward (Chapter 3), preventing inappropriate responses through its projection to the Acb (Parkinson *et al.*, 2000c). Ventral striatal DA may enhance ongoing responding, but may also 'teach' the striatum; it is speculated that this can lead to a permanent enhancement in the motivational impact or salience of a CS, or its ability to induce certain patterns of motivated response (cf. Robinson & Berridge, 1993) (and see text, pp. 243/245).

corticostriatal projections are glutamatergic and excitatory. Within the major corticostriatal loops (skeletal motor, oculomotor, 'cognitive', and limbic), there are parallel channels: circuits that maintain a degree of functional segregation (Alexander *et al.*, 1986) and that may compete for output (Deniau *et al.*, 1982). Striatal output circuitry may operate on a disinhibitory principle: GABAergic neurons in the globus pallidus and SNr tonically inhibit thalamocortical circuits, and activity in GABAergic striatal neurons can inhibit globus pallidus/SNr neurons, disinhibiting the cortex (see Alexander & Crutcher, 1990; Chevalier & Deniau, 1990). This disinhibition does not itself trigger behaviour, but *permits* it (reviewed by Chevalier & Deniau, 1990), as the striatum does not generate simple behaviour patterns, but chooses and/or links them. This concept has been well illustrated by studies of grooming in rats; small ($\sim 1 \text{ mm}^3$) excitotoxic lesions of the dorsal striatum can impair the *sequence* of grooming behaviour without affecting the rat's capacity to emit any component of the sequence (see e.g. Aldridge *et al.*, 1993; Cromwell & Berridge, 1996). Those studies that have explicitly looked at switching are also consistent with this hypothesis; thus Acb lesions have been shown to impair 'strategy switching' in a reversal situation (Reading & Dunnett, 1991), though not on all tasks (Stern & Passingham, 1995). Chapter 4 provided further evidence for a role of the ventral striatum in the direction of ongoing behaviour by conditioned stimuli, distinguishing in addition between the AcbSh, which provided the 'vigour' for Pavlovian-instrumental transfer, and the AcbC, which provided the direction or response specificity. The manner in which the AcbC and AcbSh interact in PIT is not yet clear, as for CRf (see Chapter 1), but the 'vigour'/'direction' hypothesis is consistent with theories postulating a hierarchy even within the ventral striatum, from shell to core (e.g. Haber *et al.*, 2000).

2. Acute modulation of striatal function by dopamine

Whilst glutamatergic afferents to the striatum constitute high-bandwidth pathways, capable of carrying a large amount of information, the dopaminergic input is a low-bandwidth pathway (Schultz, 1994; Mirenowicz & Schultz, 1996; Zoli *et al.*, 1998), consistent with a role in modulating other information passing through the striatum (the functions of striatal dopamine are of necessity entirely constrained by the underlying function of the striatum). Direct evidence for such a modulatory role is provided by the conditioned reinforcement paradigm (Taylor & Robbins, 1984): infusion of dopaminergic agonists into the Acb increases the rate (i.e. the momentary probability) of responding for a conditioned reinforcer, but can only 'amplify' this effect of conditioned reinforcers when information about them is arriving via glutamatergic afferents, in this case from the BLA (Cador *et al.*, 1989; Burns *et al.*, 1993).

At a cellular level in the striatum, dopamine probably focuses activity by increasing output from the most active medium spiny neurons (which are in the minority) and decreasing output from the less active cells (see Grace, 1987; Wickens & Kötter, 1995). This is mirrored at a behavioural level; increasing doses of dopamine agonists produce higher rates of activity in more and more limited categories of response (Lyon & Robbins, 1975) until stereotypy ensues.

The differences in the functions of dopamine in the dorsal and ventral striatum (reviewed by Robbins & Everitt, 1992) can then be viewed as a common action of dopamine on striatal circuits that switch different aspects of behaviour (cf. Alexander *et al.*, 1986). In the dorsal striatum, dopaminergic agonists alter the relative probability of simple motor acts, leading to stereotypy at high doses. Antagonists and dopamine depletion prevent relevant stimuli from eliciting simple motor responses, including consummatory responses; the spectrum is from a slowed response to akinesia. Similarly, dopamine depletion of the dorsal striatum impairs learning and performance of tasks based on stimulus–response decision rules (Robbins *et al.*, 1990a). As would be predicted from the corticostriatal loop account, cognitive aspects of stimulus–response coupling, such as the establishment and maintenance of an attentional or response 'set', are probably also impaired by dorsal striatal dopamine depletion (see Marsden, 1992; Robbins & Everitt, 1992, p. 122). This description emphasizes the role of the striatum as a device that selects behavioural output *in appropriate stimulus situations*.

In the ventral striatum, dopamine agonists and antagonists similarly increase or decrease the probability of stimuli affecting ongoing behaviour, but the behaviour so altered is qualitatively different. When intra-Acb amphetamine is given to rats responding for CRf, the response that is potentiated is a complex motor act, arbitrarily chosen by the experimenter, and induced by a process of conditioned motivation. The ventral striatum also mediates motivational influences on locomotion and on preparatory aspects of behaviour (Robbins & Everitt, 1992). Switching between complex behaviours is itself reduced by dopamine depletion or antagonist injection into the Acb (Koob *et al.*, 1978; Robbins & Koob, 1980; Evenden & Carli, 1985; Bakshi & Kelley, 1991).

3. The striatum and learning

The question of whether the striatum itself is involved in learning is controversial. If the switching hypothesis is correct, then striatal learning would manifest itself as a permanent change in the probability of a particular cortical pattern or behaviour being disinhibited by the striatum, given a certain pattern of inputs. Such a mechanism would also be capable of learning motor sequences. As discussed in Chapter 1, a role for the basal ganglia in habit formation was originally suggested by Mishkin *et al.* (1984), who saw a habit as a direct stimulus–response association that was learned slowly but was stable. Much of the subsequent work on this issue has proved controversial (see Wise, 1996; Wise *et al.*, 1996; White, 1997),

though Packard & McGaugh (1996) have provided good evidence for a long term change in behaviour that is dependent on the striatum. In their experiment, described in Chapter 1 (p. 46), rats were trained in a T-maze with one arm consistently baited. This task is soluble by repeating the reinforced response, or by approaching the place where food was found (a 'place response'), and these alternatives were distinguished by letting the rat approach the choice point from the opposite direction. After 8 days of training, most rats made place responses, which depended on the function of the dorsal hippocampus but not of the dorsolateral caudate nucleus. After 16 days of training, however, most rats instead made the motor response that had been reinforced. Inactivating the caudate with lidocaine eliminated this tendency and reinstated place responding, whilst inactivation of the hippocampus had no effect. Therefore, in this task, development of a stimulus-to-motor response mapping takes place slowly during reinforced training and comes to dominate behaviour, and its performance depends on the caudate nucleus. However, this response has not yet been characterized as a habit by reinforcer devaluation techniques; similarly, it is not clear from this type of experiment whether the caudate itself is the critical site of plasticity or is merely involved in behavioural expression of the response.

4. Dopaminergic effects on striatal learning; implications for addiction

Dopamine has been widely suggested to affect learning by effects exerted within the striatum. At a cellular level, dopamine can mediate heterosynaptic plasticity in the striatum (reviewed by Wickens & Köster, 1995): pre- and postsynaptic activity in the corticostriatal pathway produces long-term depression (Calabresi *et al.*, 1992) but phasic dopamine may reverse this, producing a potentiation (Wickens *et al.*, 1996) (though see Pennartz *et al.*, 1993). Single-cell recording has shown that dopaminergic neurons of the SNc/VTA respond to unpredicted rewards; with training, this response transfers to stimuli predictive of rewards (Schultz *et al.*, 1993; Mirenowicz & Schultz, 1994; Mirenowicz & Schultz, 1996). Based on the response properties of midbrain dopamine neurons, computational neuroscientists have suggested that by signalling reward prediction errors, dopamine acts as a teaching signal for striatal learning (Houk *et al.*, 1995; Montague *et al.*, 1996; Schultz *et al.*, 1997) in a system based upon temporal difference (TD) learning (Sutton, 1988), with dopamine increasing the probability of repeating responses that lead to reward. It would certainly be maladaptive to develop inflexible, habitual behaviour if such learning were not guided by a signal at least correlated with reinforcement, and the dopamine signal fulfils this property.

While the suggestion that dopamine acts as a teaching signal is controversial (e.g. Pennartz, 1995; Redgrave *et al.*, 1999b) — for example, many effects of dopaminergic manipulations are interpretable as effects on attentional or response switching — there is some behavioural evidence for dopaminergic consolidation of S–R learning. The 'win-stay' radial maze task may be solved by a stimulus–response rule, as approach to an illuminated arm is always rewarded. Performance on this task is also sensitive to caudate lesions (Packard *et al.*, 1989) and improved by post-training injections of dopamine agonists into the caudate (Packard & White, 1991). These effects are neurally and behaviourally specific: caudate manipulations had no effect on a 'win-shift' task in the same apparatus, and were doubly dissociated from the effects of lesions of the hippocampus or post-training hippocampal injections of dopaminergic agonists. Post-training microinjections represent a critical experimental test for the demonstration of task consolidation, as they cannot affect task performance. However, the task cannot be characterized as a stimulus–response habit as clearly as the T-maze task.

Does ventral striatal dopamine consolidate learning of a stimulus–motivation mapping in a similar manner? At present, this is an unanswered question. Unpublished observations from our laboratory indicate that rats responding for a conditioned reinforcer in extinction under saline conditions respond more if they have previously responded with intra-accumbens amphetamine, which is contrary to the general ten-

dency for responding to extinguish (R.N. Cardinal, T.W. Robbins and B.J. Everitt, unpublished observations). However, these data are confounded by response generalization effects (the possibility that the rats responded more simply because they had a history of high responding in the same environment). Post-training injections of the dopamine D₂ antagonist sulpiride into the Acb have been shown to impair water-maze performance (Setlow & McGaugh, 1998; 1999), but the theoretical basis of this task is not clear.

Since the dorsal striatum is involved in the development of stimulus–response habits (Reading *et al.*, 1991; Packard & McGaugh, 1996), whilst the ventral striatum is involved in motivational processes (Robbins & Everitt, 1992), a qualitative difference may exist between the two. However, if the two structures (at least, the dorsal striatum and the AcbC) perform similar functions at a neural level, then a direct comparison may be fruitful. An S–R habit may be defined as the production of a motor response with a fixed probability given a set of stimuli; that is, a simple and inflexible input/output mapping. Habits are also learned slowly. But if the striatum subserves S–R habits, then the stimulus is whatever cortical inputs arrive at a striatal segment, which depends upon the corticostriatal loop of which it is part, and the response is the pattern that the striatum consequently induces in the structures to which it projects. For the ventral striatum, the equivalent habit would be the *inflexible generation of a motivational effect* in a particular context.

Such ‘motivational habits’ may be of critical importance in the phenomenon of drug addiction (Figure 102). Compulsive drug use is characterized by behaviour that is inflexible, for it persists despite considerable cost to the addict and may become dissociated from subjective measures of drug value (Robinson & Berridge, 1993), may be elicited by specific environmental stimuli (O’Brien *et al.*, 1986), and yet involves complex, goal-directed behaviours for procuring and self-administering a drug. In a behavioural hierarchy, inappropriate reinforcement of low-level behaviours may be of trivial consequence whereas drug-induced reinforcement of a motivational process that has flexible cognitive and motor systems at its disposal may be far more destructive.

A critical question regarding the neural basis of addiction is what differentiates the effects of abused drugs from the effects of natural reinforcers, and whether this is a qualitative or a quantitative difference. In addition to the ability of drugs of abuse to activate DA systems more consistently than food reinforcers (see Di Chiara, 1998), recent evidence suggests that one effect unique to such drugs may be *sensitization*, the phenomenon by which repeated drug administration leads to an enhanced response to the drug (for reviews, see Robinson & Berridge, 1993; Altman *et al.*, 1996, pp. 302–304). Sensitization to amphetamine is induced via the drug’s effects on VTA cell bodies and is expressed as hypersensitivity to amphetamine at dopaminergic terminals in the Acb (Cador *et al.*, 1995; see also Stephens, 1995; Kalivas *et al.*, 1998; Mead & Stephens, 1998). The Pavlovian motivational process suggested to be subserved by ventral striatal dopamine (as discussed in Chapter 1, p. 46) and termed *incentive salience* or ‘wanting’ by Robinson & Berridge (Robinson & Berridge, 1993; Berridge & Robinson, 1998) has been suggested to sensitize (Robinson & Berridge, 1993); ‘incentive sensitization’ may be an important contributor to addiction. The potential link to PIT and amphetamine potentiation of CRf is clear. PIT may be an important basis for conditioned reinforcement (see Chapters 1 & 4, pp. 31/50/146) and intra-accumbens amphetamine potentiates PIT (Wyvell & Berridge, 2000, using a food US). Sensitization interacts with these Acb-dependent processes: amphetamine sensitization leads to enhanced conditioned approach and conditioned increases in amygdala dopamine in response to a CS (Harmer & Phillips, 1999), while repeated cocaine administration sensitizes the response to intra-accumbens amphetamine when responding for CRf (Taylor & Horger, 1999, using a water US). It seems likely that PIT would sensitize in a similar way; this will be an important suggestion to test, and if confirmed it will be particularly interesting to test whether psy-

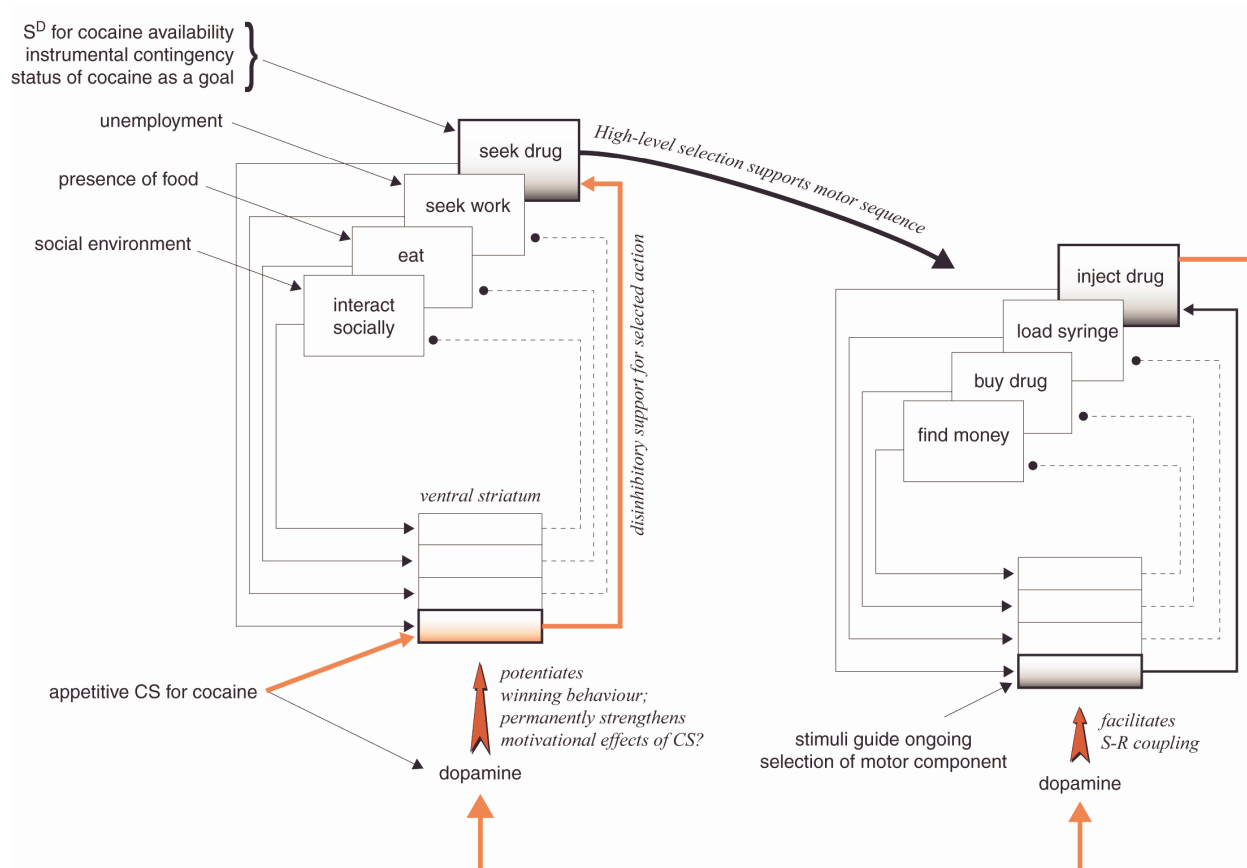


Figure 102. Highly speculative version of Figure 100 illustrating in red the particular problem posed by drugs of abuse, such as cocaine. While goal-directed behaviour may lead an individual to take such drugs (top left), just as it leads to other goals, drugs of abuse are particularly powerful at activating dopamine systems. It is possible that ventral striatal DA permanently enhances the motivational impact of Pavlovian CSs (bottom left), making those CSs potent at influencing instrumental behaviour (cf. Robinson & Berridge, 1993). If this DA system were abnormally enhanced, the Pavlovian CS might become capable of triggering complex drug-seeking behaviour even if the drug did not have high instrumental incentive value — a motivational habit.

chostimulant sensitization enhances PIT regardless of the US (for example, whether repeated noncontingent amphetamine would enhance PIT using a food US), as suggested by the comparison between Taylor & Horger (1999) and Wyvell & Berridge (2000), or whether such sensitization would predominantly affect PIT for psychostimulant USs. PIT may be especially important in addiction (with potential roles in acquisition, maintenance, and cue-induced relapse; see e.g. Tiffany & Drobles, 1990; Gawin, 1991; O'Brien *et al.*, 1998) as it represents a mechanism by which uncontrolled (noncontingent) stimuli can radically affect goal-directed responding.

5. Incorporation of the present findings relating to delayed reinforcement

The finding that the AcbC is critical for choosing delayed reward is intriguing and novel, but the psychological basis of this effect is not yet clear. As discussed in Chapter 7 (p. 226), it will be important to establish whether this deficit is entirely due to a difference in the perception of reward magnitude in AcbC-lesioned rats, or whether a specific delay-dependent deficit exists. Furthermore, as discussed above (p. 239), the psychological processes contributing to choice at different delays are not well understood at present. It is therefore unclear whether the deficit in AcbC-lesioned rats can be interpreted entirely within the framework of ventral striatal function reviewed here and by Parkinson *et al.* (2000a). If it can be shown behaviourally that Pavlovian conditioned motivation is a major contributor to preference for delayed rewards, this may be accomplished. However, as the task used in Chapter 7 had no explicit cues signalling

the delayed reward, this explanation appears specious and *post hoc*, and theories of Acb function may have to be extended (as discussed in Chapter 7, p. 230) to accommodate the novel finding.

It will therefore be important to assess the roles of the AcbSh and Acb DA in preference for delayed reward, using excitotoxic and DA-depleting lesions respectively, particularly given the evidence from studies of ADHD and animal models thereof implicating ventral striatal dopamine in the pathogenesis of impulsive choice (reviewed in Chapters 6 & 7).

To examine the possibility that the AcbC plays a wider role in learning across delays, it will be interesting to establish whether AcbC-lesioned rats are impaired at instrumental learning when the reinforcer is delayed (in the absence of an immediately-available alternative). Demonstration of such an impairment would tend to support the view that AcbC-lesioned rats prefer immediate reinforcement because they have difficulty learning that the alternative choice leads to reinforcement at all, while failure to find an impairment would suggest that AcbC-lesioned rats are 'aware of their options' even as they choose the immediate reinforcer. More generally still, it remains to be established whether AcbC lesions impair *Pavlovian* conditioning when the CS–US interval is long (trace conditioning). Additionally, trace conditioning may be less effective than conditioning with a short CS–US interval because trace conditioning promotes conditioning to other stimuli occurring during the interval, including contextual stimuli (see Dickinson, 1980, pp. 61–70; Mackintosh, 1983, pp. 202–210). AcbC lesions have been shown to impair conditioning to discrete cues but enhance contextual conditioning in a lick suppression task (Parkinson *et al.*, 1999c). The possibility may be entertained that hippocampal lesions (which disrupt contextual conditioning in the same task; Selden *et al.*, 1991) might promote responding for delayed reinforcement by reducing contextual overshadowing in instrumental learning.

Finally, even if AcbC lesions are shown to cause a purely delay-dependent (rather than reward magnitude-dependent) deficit, it would be extremely unusual for striatal lesions to impair a behaviour not impaired by lesions to its afferents (other than behavioural sequencing and switching, as discussed above, pp. 241–243). Thus, before ascribing a specific delay-dependent function to the Acb, it must be shown that lesions of the afferents to the Acb do not produce the same deficit. This work was begun in the present thesis with the demonstration that ACC and mPFC lesions do not impair rats' ability to choose a delayed reward, but the effect of lesions to other glutamatergic afferents such as the BLA, the subiculum, and the orbitofrontal cortex are unknown, as are the effects of manipulations of the DA and 5-HT innervation of the Acb. The task developed by Evenden & Ryan (1996) has proved very useful in this field of study. It has now been successfully applied to pharmacological, behavioural, and lesion studies of delayed reinforcement, and will likely prove a good starting point for future work to elucidate further the neural circuit responsible for the important ability of animals to gain reinforcement, even when it is delayed.

Reinforcement learning in the brain: an integrative view

Reinforcement is not unitary. As reviewed in Chapter 1, Pavlovian conditioning creates multiple representations. Their neural bases are gradually becoming clear. These include CS–US(sensory) or S–S associations, required for sensory preconditioning and dependent at least in part on the perirhinal cortex for visual stimuli and on the gustatory neocortex for food USs; CS–US(motivational) associations, suggested to depend on the BLA; direct CS–affect associations, which are responsible for transreinforcer blocking and are poorly understood; and CS–response associations, whose neural basis depends on the specific response (being cerebellum-dependent in the case of discrete skeletomotor CRs, and CeA-dependent in the case of several others such as conditioned suppression and PIT). Learning theorists discovered the

existence of these multiple representations and learning theory has dramatically enhanced neurobiological studies of conditioning, but sometimes neural dissociations within conditioning have been found that were not predicted by learning theory (e.g. Steinmetz, 2000); in these situations, neurobiology can inform learning theory. When considering the manner in which representations change with training, and how these representations are formed and interact across widely distributed neural systems, neither behavioural studies nor biology have provided clear answers and much work remains to be done.

Other structures contribute to instrumental conditioning (also reviewed in Chapter 1), which also creates multiple representations and can be heavily influenced by Pavlovian conditioning procedures. The prefrontal (prelimbic) cortex is critical for the perception of instrumental contingencies, while gustatory neocortex also has a role in recalling the instrumental incentive values of foodstuffs. It is not yet known how either structure acquires or represents this information, or how they interact with other representations of stimulus and reward value such as those in the amygdala and orbitofrontal cortex. It seems likely that the dorsal striatum contributes in some way to the acquisition of S–R responding, but this requires definitive proof. The nucleus accumbens was accurately described by Mogenson *et al.* (1980) as a limbic–motor interface, but may also be considered a Pavlovian–instrumental interface; it is a critical site for the motivational and directional impact of Pavlovian CSs on instrumental responding and locomotor approach. This multiplicity of representations should guide modelling studies: simple computational models of reinforcement learning, typically S–R in nature, may provide useful information regarding the principles upon which S–R systems can operate, but are often inadequate for describing simple instrumental behaviour in rats. At least some of the processes governing instrumental responding are based on declarative knowledge that is akin to symbolic processing, and yet these complex representations are known to interact with each other and with basic motivational states. Understanding this interface, and with it the nature of neural representations themselves, is one of the greatest challenges for neurobiology.

This thesis has not answered or even addressed the vast majority of these questions, but it has provided evidence that the anterior cingulate cortex makes a discriminative contribution to Pavlovian conditioning; it has elucidated further the manner in which the nucleus accumbens core and shell mediate the impact of Pavlovian CSs on instrumental responding; it has demonstrated an interaction between Pavlovian CSs and the effects of psychostimulant drugs on choice of delayed reinforcement, and it has demonstrated that the nucleus accumbens core is a critical part of the neural circuitry mediating the effects of delayed reinforcement on instrumental responding.

References

- Aberman, J. E. & Salamone, J. D. (1999). Nucleus accumbens dopamine depletions make rats more sensitive to high ratio requirements but do not impair primary food reinforcement. *Neuroscience* **92**: 545-552.
- Adams, C. D. (1982). Variations in the sensitivity of instrumental responding to reinforcer devaluation. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **34**: 77-98.
- Adams, C. D. & Dickinson, A. (1981). Instrumental responding following reinforcer devaluation. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **33**: 109-121.
- Ainslie, G. (1975). Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological Bulletin* **82**: 463-496.
- Alderson, H. L., Robbins, T. W. & Everitt, B. J. (in press-a). The effects of excitotoxic lesions of the basolateral amygdala on the acquisition of heroin-seeking behaviour in rats [DOI 10.1007/s002130000527]. *Psychopharmacology*.
- Alderson, H. L., Robbins, T. W. & Everitt, B. J. (in press-b). Heroin self-administration under a second-order schedule of reinforcement: acquisition and maintenance of heroin-seeking behaviour in rats [DOI 10.1007/s002130000429]. *Psychopharmacology*.
- Aldridge, J. W., Berridge, K. C., Herman, M. & Zimmer, L. (1993). Neuronal coding of serial order: syntax of grooming in the neostriatum. *Psychological Science* **4**: 391-395.
- Alexander, G. E. & Crutcher, M. D. (1990). Functional architecture of basal ganglia circuits: neural substrates of parallel processing. *Trends in Neurosciences* **13**: 266-271.
- Alexander, G. E., Crutcher, M. D. & DeLong, M. R. (1990). Basal ganglia-thalamocortical circuits - parallel substrates for motor, oculomotor, prefrontal and limbic functions. *Progress in Brain Research* **85**: 119-146.
- Alexander, G. E., DeLong, M. R. & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience* **9**: 357-381.
- Allport, A. (1987). Selection for action: Some behavioral and neurophysiological considerations of attention and action. In *Perspectives on Perception and Action* (Heuer, H. & Sanders, A. F., eds.), pp. 395-419. Lawrence Erlbaum Associates, Hillsdale, New Jersey.
- Al-Ruwaitea, A. S., Al-Zahrani, S. S., Ho, M. Y., Bradshaw, C. M. & Szabadi, E. (1997a). Effect of central 5-hydroxytryptamine depletion on performance in the "time-left" procedure: further evidence for a role of the 5-hydroxytryptaminergic pathways in behavioural "switching". *Psychopharmacology* **134**: 179-186.
- Al-Ruwaitea, A. S. A., Al-Zahrani, S. S. A., Ho, M. Y., Bradshaw, C. M. & Szabadi, E. (1997b). 5-Hydroxytryptamine and interval timing. In *Time and Behaviour: Psychological and Neurobehavioural Analyses* (Bradshaw, C. M. & Szabadi, E., eds.), pp. 517-570. Elsevier, Amsterdam.
- Altman, J., Everitt, B. J., Glautier, S., Markou, A., Nutt, D., Oretti, R., Phillips, G. D. & Robbins, T. W. (1996). The biological, social and clinical bases of drug addiction: commentary and debate. *Psychopharmacology* **125**: 285-345.
- Al-Zahrani, S. S., Ho, M. Y., Velazquez Martinez, D. N., Lopez Cabrera, M., Bradshaw, C. M. & Szabadi, E. (1996). Effect of destruction of the 5-hydroxytryptaminergic pathways on behavioural timing and "switching" in a free-operant psychophysical procedure. *Psychopharmacology* **127**: 346-352.
- Al-Zahrani, S. S. A., Ho, M. Y., Al-Ruwaitea, A. S. A., Bradshaw, C. M. & Szabadi, E. (1997). Effect of destruction of the 5-hydroxytryptaminergic pathways on temporal memory: Quantitative analysis with a delayed interval bisection task. *Psychopharmacology* **129**: 48-55.
- Amaral, D. G., Price, J. L., Pitkänen, A. & Carmichael, S. T. (1992). Anatomical organization of the primate amygdaloid complex. In *The amygdala: neurobiological aspects of emotion, memory, and mental dysfunction* (Aggleton, J. P., ed.), pp. 1-66. Wiley-Liss, New York.
- Anand, B. K. & Brobeck, J. R. (1951). Hypothalamic control of food intake. *Yale Journal of Biology and Medicine* **24**: 123-140.
- APA (1994). *Diagnostic and Statistical Manual of Mental Disorders, version IV (DSM-IV)*, American Psychiatric Association, Washington DC.
- Aristotle (350 BC / 1925). *Nicomachean Ethics [translated by W.D. Ross]*, Clarendon Press, Oxford.

- Arroyo, M., Markou, A., Robbins, T. W. & Everitt, B. J. (1998). Acquisition, maintenance and reinstatement of intravenous cocaine self-administration under a second-order schedule of reinforcement in rats: effects of conditioned cues and continuous access to cocaine. *Psychopharmacology* **140**: 331-344.
- Åsberg, M., Träskman, L. & Thorén, P. (1976). 5-HIAA in the cerebrospinal fluid: a biochemical suicide predictor. *Archives of General Psychiatry* **33**: 1193-1197.
- Autor, S. M. (1969). The strength of conditioned reinforcers as a function of frequency and probability of reinforcement. In *Conditioned reinforcement* (Hendry, D. P., ed.), pp. 127-162. The Dorsey Press, Homewood, Illinois.
- Awh, E. & Gehring, W. J. (1999). The anterior cingulate cortex lends a hand in response selection. *Nature Neuroscience* **2**: 853-854.
- Baars, B. J. (1988). *A cognitive theory of consciousness*, Cambridge University Press, Cambridge.
- Baer, L., Rauch, S. L., Ballantine, H. T., Jr., Martuza, R., Cosgrove, R., Cassem, E., Giriunas, I., Manzo, P. A., Dimino, C. & Jenike, M. A. (1995). Cingulotomy for intractable obsessive-compulsive disorder. Prospective long-term follow-up of 18 patients. *Archives of General Psychiatry* **52**: 384-392.
- Baeyens, F., Eelen, P., van den Berg, H. & Crombez, G. (1990). Flavor-flavor and color-flavor conditioning in humans. *Learning & Motivation* **21**: 434-455.
- Bakshi, V. P. & Kelley, A. E. (1991). Dopaminergic regulation of feeding behavior. 1: Differential effects of haloperidol microinfusion into 3 striatal subregions. *Psychobiology* **19**: 223-232.
- Balleine, B. (1992). Instrumental performance following a shift in primary motivation depends on incentive learning. *Journal of Experimental Psychology: Animal Behavior Processes* **18**: 236-250.
- Balleine, B. (1994). Asymmetrical interactions between thirst and hunger in Pavlovian-instrumental transfer. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **47**: 211-231.
- Balleine, B., Ball, J. & Dickinson, A. (1994). Benzodiazepine-induced outcome revaluation and the motivational control of instrumental action in rats. *Behavioral Neuroscience* **108**: 573-589.
- Balleine, B., Davies, A. & Dickinson, A. (1995a). Cholecystokinin attenuates incentive learning in rats. *Behavioral Neuroscience* **109**: 312-319.
- Balleine, B. & Dickinson, A. (1991). Instrumental performance following reinforcer devaluation depends upon incentive learning. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **43**: 279-296.
- Balleine, B. & Dickinson, A. (1992). Signalling and incentive processes in instrumental reinforcer devaluation. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **45**: 285-301.
- Balleine, B. & Dickinson, A. (1994). Role of cholecystokinin in the motivational control of instrumental action in rats. *Behavioral Neuroscience* **108**: 590-605.
- Balleine, B. & Killcross, S. (1994). Effects of ibotenic acid lesions of the nucleus accumbens on instrumental action. *Behavioural Brain Research* **65**: 181-193.
- Balleine, B. W. & Dickinson, A. (1998a). Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology* **37**: 407-419.
- Balleine, B. W. & Dickinson, A. (1998b). The role of incentive learning in instrumental outcome revaluation by sensory-specific satiety. *Animal Learning & Behavior* **26**: 46-59.
- Balleine, B. W. & Dickinson, A. (2000). The effect of lesions of the insular cortex on instrumental conditioning: evidence for a role in incentive memory. *Journal of Neuroscience* **20**: 8954-8964.
- Balleine, B. W., Garner, C., Gonzalez, F. & Dickinson, A. (1995b). Motivational control of heterogeneous instrumental chains. *Journal of Experimental Psychology: Animal Behavior Processes* **21**: 203-217.
- Barratt, E. S. & Patton, J. H. (1983). Impulsivity: cognitive, behavioral and psychophysiological correlates. In *Biological bases of sensation seeking, impulsivity, and anxiety* (Zuckerman, M., ed.), pp. 77-116. Lawrence Erlbaum Associates, Hillsdale, New Jersey.
- Barto, A. G. (1995). Adaptive critics and the basal ganglia. In *Models of Information Processing in the Basal Ganglia* (Houk, J. C., Davis, J. L. & Beiser, D. G., eds.), pp. 215-232. MIT Press, Cambridge, Massachusetts.

- Bassareo, V. & Di Chiara, G. (1999). Differential responsiveness of dopamine transmission to food-stimuli in nucleus accumbens shell/core compartments. *Neuroscience* **89**: 637-641.
- Baum, W. M. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior* **22**: 231-242.
- Baum, W. M. (1979). Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior* **32**: 269-281.
- Baum, W. M. & Rachlin, H. C. (1969). Choice as time allocation. *Journal of the Experimental Analysis of Behavior* **12**: 861-874.
- Bautista, L. M., Alonso, J. C. & Alonso, J. A. (1992). A 20-year study of wintering common crane fluctuations using time series analysis. *Journal of Wildlife Management* **56**: 563-572.
- Baxter, D. J. & Zamble, E. (1982). Reinforcer and response specificity in appetitive transfer of control. *Animal Learning & Behavior* **10**: 201-210.
- Baxter, M. G., Parker, A., Lindner, C. C. C., Izquierdo, A. D. & Murray, E. A. (2000). Control of response selection by reinforcer value requires interaction of amygdala and orbital prefrontal cortex. *Journal of Neuroscience* **20**: 4311-4319.
- Bechara, A., Damasio, H. & Damasio, A. R. (2000). Emotion, decision making and the orbitofrontal cortex. *Cerebral Cortex* **10**: 295-307.
- Bechara, A., Damasio, H., Damasio, A. R. & Lee, G. P. (1999). Different contributions of the human amygdala and ventromedial prefrontal cortex to decision-making. *Journal of Neuroscience* **19**: 5473-5481.
- Bechara, A., Damasio, H., Tranel, D. & Anderson, S. W. (1998). Dissociation of working memory from decision making within the human prefrontal cortex. *Journal of Neuroscience* **18**: 428-437.
- Bechara, A., Damasio, H., Tranel, D. & Damasio, A. R. (1997). Deciding advantageously before knowing the advantageous strategy. *Science* **275**: 1293-1295.
- Bechara, A., Tranel, D., Damasio, H. & Damasio, A. R. (1996). Failure to respond autonomically to anticipated future outcomes following damage to prefrontal cortex. *Cerebral Cortex* **6**: 215-225.
- Belke, T. W. (1998). Qualitatively different reinforcers and parameters of Herrnstein's (1970) response-strength equation. *Animal Learning & Behavior* **26**: 235-242.
- Bench, C. J., Friston, K. J., Brown, R. G., Scott, L. C., Frackowiak, R. S. & Dolan, R. J. (1992). The anatomy of melancholia - focal abnormalities of cerebral blood flow in major depression. *Psychological Medicine* **22**: 607-615.
- Bench, C. J., Frith, C. D., Grasby, P. M., Friston, K. J., Paulesu, E., Frackowiak, R. S. & Dolan, R. J. (1993). Investigations of the functional anatomy of attention using the Stroop test. *Neuropsychologia* **31**: 907-922.
- Berendse, H. W., Galisdegraaf, Y. & Groenewegen, H. J. (1992). Topographical organization and relationship with ventral striatal compartments of prefrontal corticostriatal projections in the rat. *Journal of Comparative Neurology* **316**: 314-347.
- Berns, G. S., Cohen, J. D. & Mintun, M. A. (1997). Brain regions responsive to novelty in the absence of awareness. *Science* **276**: 1272-1275.
- Bernstein, P. S., Scheffers, M. K. & Coles, M. G. H. (1995). Where did I go wrong - a psychophysiological analysis of error-detection. *Journal of Experimental Psychology: Human Perception and Performance* **21**: 1312-1322.
- Berridge, K. C. (1991). Modulation of taste affect by hunger, caloric satiety, and sensory-specific satiety in the rat. *Appetite* **16**: 103-120.
- Berridge, K. C. (1996). Food reward: Brain substrates of wanting and liking. *Neuroscience and Biobehavioral Reviews* **20**: 1-25.
- Berridge, K. C., Flynn, F. W., Schulkin, J. & Grill, H. J. (1984). Sodium depletion enhances salt palatability in rats. *Behavioral Neuroscience* **98**: 652-660.
- Berridge, K. C., Grill, H. J. & Norgren, R. (1981). Relation of consummatory responses and preabsorptive insuline release to palatability and learned taste aversions. *Journal of Comparative and Physiological Psychology* **95**: 363-382.
- Berridge, K. C. & Robinson, T. E. (1998). What is the role of dopamine in reward: hedonic impact, reward learning, or incentive salience? *Brain Research Reviews* **28**: 309-369.
- Bickel, W. K., Odum, A. L. & Madden, G. J. (1999). Impulsivity and cigarette smoking: delay discounting in current, never, and ex-smokers. *Psychopharmacology* **146**: 447-454.

- Bizot, J., Le Bihan, C., Puech, A. J., Hamon, M. & Thiébot, M. (1999). Serotonin and tolerance to delay of reward in rats. *Psychopharmacology* **146**: 400-412.
- Bizot, J. C. & Thiébot, M. H. (1996). Impulsivity as a confounding factor in certain animal tests of cognitive function. *Cognitive Brain Research* **3**: 243-250.
- Blackburn, J. R., Phillips, A. G. & Fibiger, H. C. (1987). Dopamine and preparatory behavior. 1. Effects of pimozide. *Behavioral Neuroscience* **101**: 352-360.
- Blundell, J. E., Strupp, B. J. & Latham, C. J. (1977). Pharmacological manipulation of hoarding: further analysis of amphetamine isomers and pimozide. *Physiological Psychology* **5**: 462-468.
- Blundell, P. J. & Killcross, A. S. (2000a). The basolateral amygdala (BLA) mediates the sensory properties of reinforcers in appetitive conditioning. *European Journal of Neuroscience* **12 (supplement 11)**: 78.
- Blundell, P. J. & Killcross, A. S. (2000b). Effects of excitotoxic lesions of the basolateral nucleus of the amygdala on associative learning in rats. *Journal of Psychopharmacology* **14 (3) (supplement)**: A48.
- Bolles, R. C. & Fanselow, M. S. (1980). A perceptual-defensive-recuperative model of fear and pain. *Behavioral Brain Sciences* **3**: 291-323.
- Bolles, R. C., Holtz, R., Dunn, T. & Hill, W. (1980). Comparison of stimulus learning and response learning in a punishment situation. *Learning & Motivation* **11**: 78-96.
- Botvinick, M., Nystrom, L. E., Fissell, K., Carter, C. S. & Cohen, J. D. (1999). Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature* **402**: 179-181.
- Box, G. E. P. (1954). Some theorems on quadratic forms applied in the study of analysis of variance problems: I. Effect of inequality of variance in the one-way classification. *Annals of Mathematical Statistics* **25**: 290-302.
- Box, G. E. P. & Jenkins, G. M. (1970). *Time series analysis: forecasting and control*, Holden-Day, San Francisco.
- Bradley, C. (1937). The behavior of children receiving Benzedrine. *American Journal of Psychiatry* **94**: 577-585.
- Bradley, D. R. & Russell, R. L. (1998). Some cautions regarding statistical power in split-plot designs. *Behavior Research Methods, Instruments and Computers* **30**: 462-477.
- Bradshaw, C. M. & Szabadi, E. (1992). Choice between delayed reinforcers in a discrete-trials schedule - the effect of deprivation level. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **44B**: 1-16.
- Branch, M. N. (1975). Effects of chlorpromazine and d-amphetamine on observing responses during a fixed-interval schedule. *Psychopharmacologia* **42**: 87-93.
- Breiter, H. C., Rauch, S. L., Kwong, K. K., Baker, J. R., Weisskoff, R. M., Kennedy, D. N., Kendrick, A. D., Davis, T. L., Jiang, A. P., Cohen, M. S., Stern, C. E., Belliveau, J. W., Baer, L., Osullivan, R. L., Savage, C. R., Jenike, M. A. & Rosen, B. R. (1996). Functional magnetic resonance imaging of symptom provocation in obsessive-compulsive disorder. *Archives of General Psychiatry* **53**: 595-606.
- Broca, P. (1878). Anatomie comparée des circonvolutions cérébrales. Le grand lobe limbique et la scissure limbique dans la série des mammifères. *Revue Anthropologique (Paris)* **1**: 456-498.
- Brog, J. S., Salyapongse, A., Deutch, A. Y. & Zahm, D. S. (1993). The patterns of afferent innervation of the core and shell in the "accumbens" part of the rat ventral striatum: immunohistochemical detection of retrogradely transported fluoro-gold. *Journal of Comparative Neurology* **338**: 255-278.
- Brogden, W. J. (1939). Sensory pre-conditioning. *Journal of Experimental Psychology* **25**: 323-332.
- Brown, G. L. & Linnoila, M. (1990). CSF serotonin metabolite (5HIAA) studies in depression, impulsivity and violence. *Journal of Clinical Psychiatry* **51 (supplement 4)**: 31-41.
- Brown, K. S. (1999). Oops... sorry. Researchers know the brain somehow spots mistakes. *New Scientist* **13 February 1999**: 38-41.
- Brown, P. L. & Jenkins, H. M. (1968). Auto-shaping of the pigeon's keypeck. *Journal of the Experimental Analysis of Behavior* **11**: 1-8.
- Brown, V. J. & Bowman, E. M. (1995). Discriminative cues indicating reward magnitude continue to determine reaction time of rats following lesions of the nucleus accumbens. *European Journal of Neuroscience* **7**: 2479-2485.
- Browne, M. P. (1976). The role of primary reinforcement and overt movements in auto-shaping in the pigeon. *Animal Learning & Behavior* **4**: 287-292.

- Buchanan, S. L. & Powell, D. A. (1982a). Cingulate cortex: its role in Pavlovian conditioning. *Journal of Comparative and Physiological Psychology* **96**: 755-774.
- Buchanan, S. L. & Powell, D. A. (1982b). Cingulate damage attenuates conditioned bradycardia. *Neuroscience Letters* **29**: 261-268.
- Burns, L. H., Everitt, B. J. & Robbins, T. W. (1999). Effects of excitotoxic lesions of the basolateral amygdala on conditional discrimination learning with primary and conditioned reinforcement. *Behavioural Brain Research* **100**: 123-133.
- Burns, L. H., Robbins, T. W. & Everitt, B. J. (1993). Differential effects of excitotoxic lesions of the basolateral amygdala, ventral subiculum and medial prefrontal cortex on responding with conditioned reinforcement and locomotor activity potentiated by intra-accumbens infusions of D-amphetamine. *Behavioural Brain Research* **55**: 167-183.
- Busatto, G. F., Zamignani, D. R., Buchpiguel, C. A., Garrido, G. E. J., Glabus, M. F., Rocha, E. T., Maia, A. F., Rosario-Campos, M. C., Castro, C. C., Furuie, S. S., Gutierrez, M. A., McGuire, P. K. & Miguel, E. C. (2000). A voxel-based investigation of regional cerebral blood flow abnormalities in obsessive-compulsive disorder using single photon emission computed tomography (SPECT). *Psychiatry Research - Neuroimaging* **99**: 15-27.
- Bush, G., Frazier, J. A., Rauch, S. L., Seidman, L. J., Whalen, P. J., Jenike, M. A., Rosen, B. R. & Biederman, J. (1999). Anterior cingulate cortex dysfunction in attention-deficit/hyperactivity disorder revealed by fMRI and the Counting Stroop. *Biological Psychiatry* **45**: 1542-1552.
- Bush, G., Luu, P. & Posner, M. I. (2000). Cognitive and emotional influences in anterior cingulate cortex. *Trends in Cognitive Sciences* **4**: 215-222.
- Bush, G., Whalen, P. J., Rosen, B. R., Jenike, M. A., McInerney, S. C. & Rauch, S. L. (1998). The counting Stroop: an interference task specialized for functional neuroimaging - validation study with functional MRI. *Human Brain Mapping* **6**: 270-282.
- Bussey, T. J. (1996). Functions of the cingulate cortex in learning and memory in the rat. Unpublished PhD thesis, University of Cambridge.
- Bussey, T. J., Everitt, B. J. & Robbins, T. W. (1997a). Dissociable effects of cingulate and medial frontal cortex lesions on stimulus-reward learning using a novel Pavlovian autoshaping procedure for the rat: implications for the neurobiology of emotion. *Behavioral Neuroscience* **111**: 908-919.
- Bussey, T. J., Muir, J. L., Everitt, B. J. & Robbins, T. W. (1996). Dissociable effects of anterior and posterior cingulate cortex lesions on the acquisition of a conditional visual discrimination: facilitation of early learning vs. impairment of late learning. *Behavioural Brain Research* **82**: 45-56.
- Bussey, T. J., Muir, J. L., Everitt, B. J. & Robbins, T. W. (1997b). Triple dissociation of anterior cingulate, posterior cingulate, and medial frontal cortices on visual discrimination tasks using a touchscreen testing procedure for the rat. *Behavioral Neuroscience* **111**: 920-936.
- Butcher, L. L. (1995). Cholinergic neurons and networks. In *The Rat Nervous System* (Paxinos, G., ed.), pp. 1003-1015. Academic Press, London.
- Cador, M., Bjjjou, Y. & Stinus, L. (1995). Evidence of a complete independence of the neurobiological substrates for the induction and expression of behavioral sensitization to amphetamine. *Neuroscience* **65**: 385-395.
- Cador, M., Robbins, T. W. & Everitt, B. J. (1989). Involvement of the amygdala in stimulus-reward associations: interaction with the ventral striatum. *Neuroscience* **30**: 77-86.
- Cador, M., Taylor, J. R. & Robbins, T. W. (1991). Potentiation of the effects of reward-related stimuli by dopaminergic-dependent mechanisms in the nucleus accumbens. *Psychopharmacology* **104**: 377-385.
- Calabresi, P., Maj, R., Pisani, A., Mercuri, N. B. & Bernardi, G. (1992). Long-term synaptic depression in the striatum: physiological and pharmacological characterization. *Journal of Neuroscience* **12**: 4224-4233.
- Canon, J. G. (1979). A comparison of clozapine, chlorpromazine, and thioridazine upon DRL performance in the squirrel monkey. *Psychopharmacology* **64**: 55-60.
- Cardinal, R. N., Parkinson, J. A., Robbins, T. W., Dickinson, A. & Everitt, B. J. (2000a). Effects of lesions of the nucleus accumbens core and shell on response-specific Pavlovian-instrumental transfer. *Journal of Psychopharmacology* **14(3) (supplement)**: A68.
- Cardinal, R. N., Robbins, T. W. & Everitt, B. J. (2000b). The effects of d-amphetamine, chlordiazepoxide, alpha-flupenthixol and behavioural manipulations on choice of signalled and unsignalled delayed reinforcement in rats. *Psychopharmacology* **152**: 362-375.

- Carey, M. P., Diewald, L. M., Esposito, F. J., Pellicano, M. P., Carnevale, U. A. G., Sergeant, J. A., Papa, M. & Sadile, A. G. (1998). Differential distribution, affinity and plasticity of dopamine D-1 and D-2 receptors in the target sites of the mesolimbic system in an animal model of ADHD. *Behavioural Brain Research* **94**: 173-185.
- Carey, R. J. & Kritkasky, R. P. (1972). Absence of a response-rate-dependent effect of d-amphetamine on a DRL schedule when reinforcement is signaled. *Psychonomic Science* **26**: 285-286.
- Carlson, N. R. (1991). *Physiology of Behaviour*. Fourth edition, Allyn & Bacon, London.
- Carmichael, S. T. & Price, J. L. (1994). Architectonic subdivision of the orbital and medial prefrontal cortex in the macaque monkey. *Journal of Comparative Neurology* **346**: 366-402.
- Carter, C. S., Botvinick, M. M. & Cohen, J. D. (1999). The contribution of the anterior cingulate cortex to executive processes in cognition. *Reviews in the Neurosciences* **10**: 49-57.
- Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D. & Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science* **280**: 747-749.
- Carter, C. S., Macdonald, A. M., Botvinick, M., Ross, L. L., Stenger, V. A., Noll, D. & Cohen, J. D. (2000). Parsing executive processes: Strategic vs. evaluative functions of the anterior cingulate cortex. *Proceedings of the National Academy of Sciences of the United States of America* **97**: 1944-1948.
- Carter, C. S., Mintun, M. & Cohen, J. D. (1995). Interference and facilitation effects during selective attention: An (H₂O)-O-15 PET study of Stroop task performance. *Neuroimage* **2**: 264-272.
- Catania, A. C. (1970). Reinforcement schedules and psychophysical judgment: A study of some temporal properties of behavior. In *The theory of reinforcement schedules* (Schoenfeld, W. N., ed.), pp. 1-42. Appleton-Century-Crofts, New York.
- Charrier, D. & Thiébot, M. H. (1996). Effects of psychotropic drugs on rat responding in an operant paradigm involving choice between delayed reinforcers. *Pharmacology, Biochemistry and Behavior* **54**: 149-157.
- Chase, V. M., Hertwig, R. & Gigerenzer, G. (1998). Visions of rationality. *Trends in Cognitive Sciences* **2**: 206-214.
- Chevalier, G. & Deniau, J. M. (1990). Disinhibition as a basic process in the expression of striatal functions. *Trends in Neurosciences* **13**: 277-280.
- Chiba, A. A., Bucci, D. J., Holland, P. C. & Gallagher, M. (1995). Basal forebrain cholinergic lesions disrupt increments but not decrements in conditioned stimulus processing. *Journal of Neuroscience* **15**: 7315-7322.
- Childress, A. R., Mozley, P. D., McElgin, W., Fitzgerald, J., Reivich, M. & O'Brien, C. P. (1999). Limbic activation during cue-induced cocaine craving. *American Journal of Psychiatry* **156**: 11-18.
- Choi, D. W. (1988). Glutamate neurotoxicity and diseases of the nervous system. *Neuron* **1**: 623-634.
- Choi, D. W. (1995). Calcium: still centre-stage in hypoxic-ischemic neuronal death. *Trends in Neurosciences* **18**: 58-60.
- Cohen, J. D., Botvinick, M. & Carter, C. S. (2000). Anterior cingulate and prefrontal cortex: who's in control? *Nature Neuroscience* **3**: 421-423.
- Cole, B. J. & Robbins, T. W. (1989). Effects of 6-hydroxydopamine lesions of the nucleus accumbens septi on performance of a 5-choice serial reaction time task in rats: implications for theories of selective attention and arousal. *Behavioural Brain Research* **33**: 165-179.
- Coles, M. G. H., Scheffers, M. K. & Holroyd, C. (1998). Berger's dream? The error-related negativity and modern cognitive psychophysiology. In *Quantitative and topological EEG and EMG analysis* (Witte, H., Zwiener, U., Schack, B. & Döring, A., eds.), pp. 96-102. Druckhaus Mayer Verlag, Jena-Erlangen, Germany.
- Colwill, R. M. & Motzkin, D. K. (1994). Encoding of the unconditioned stimulus in Pavlovian conditioning. *Animal Learning & Behavior* **22**: 384-394.
- Colwill, R. M. & Rescorla, R. A. (1985). Postconditioning devaluation of a reinforcer affects instrumental responding. *Journal of Experimental Psychology: Animal Behavior Processes* **11**: 120-132.
- Colwill, R. M. & Rescorla, R. A. (1988). Associations between the discriminative stimulus and the reinforcer in instrumental learning. *Journal of Experimental Psychology: Animal Behavior Processes* **14**: 155-164.
- Colwill, R. M. & Rescorla, R. A. (1990). Evidence for the hierarchical structure of instrumental learning. *Animal Learning & Behavior* **18**: 71-82.

- Cools, A. R. (1980). Role of the neostriatal dopaminergic activity in sequencing and selecting behavioural strategies: facilitation of processes involved in selecting the best strategy in a stressful situation. *Behavioural Brain Research* **1**: 361-378.
- Cools, A. R., van den Bos, R., Ploeger, G. & Ellenbroek, B. A. (1991). Gating function of noradrenaline in the ventral striatum: its role in behavioural responses to environmental and pharmacological challenges. In *The Mesolimbic Dopamine System: from Motivation to Action* (Willner, P. & Scheel-Kruger, J., eds.), pp. 139-171. John Wiley and Sons, Chichester.
- Corbit, L., Muir, J. & Balleine, B. W. (submitted). The role of the nucleus accumbens in instrumental conditioning: Evidence for functional dissociation between accumbens core and shell. .
- Corbit, L. H. & Balleine, B. W. (2000a). The effects of lesions of the nucleus accumbens core and shell on tests of instrumental conditioning. *Society for Neuroscience Abstracts* **26**: 1741.
- Corbit, L. H. & Balleine, B. W. (2000b). The role of the hippocampus in instrumental conditioning. *Journal of Neuroscience* **20**: 4233-4239.
- Cousins, M. S., Atherton, A., Turner, L. & Salamone, J. D. (1996). Nucleus accumbens dopamine depletions alter relative response allocation in a T-maze cost/benefit task. *Behavioural Brain Research* **74**: 189-197.
- Cousins, M. S., Sokolowski, J. D. & Salamone, J. D. (1993). Different effects of nucleus accumbens and ventrolateral striatal dopamine depletions on instrumental response selection in the rat. *Pharmacology, Biochemistry and Behavior* **46**: 943-951.
- Coutureau, E., Dix, S. L. & Killcross, A. S. (2000). Involvement of the medial prefrontal cortex-basolateral amygdala pathway in fear-related behaviour in rats. *European Journal of Neuroscience* **12 (supplement 11)**: 156.
- Cromwell, H. C. & Berridge, K. C. (1996). Implementation of action sequences by a neostriatal site: A lesion mapping study of grooming syntax. *Journal of Neuroscience* **16**: 3444-3458.
- Damasio, A. R. (1994). *Descartes' Error*, Grosset/Putnam, New York.
- Damasio, A. R. (1996). The somatic marker hypothesis and the possible functions of the prefrontal cortex. *Philosophical Transactions of the Royal Society of London, Series B - Biological Sciences* **351**: 1413-1420.
- Damasio, A. R. & Damasio, H. (1993). Cortical systems underlying knowledge retrieval: evidence from human lesion studies. In *Exploring Brain Functions: Models in Neuroscience* (Poggio, T. A. & Glaser, D. A., eds.), pp. 233-249. John Wiley & Sons, Chichester.
- Davey, G. C., Oakley, D. A. & Cleland, G. G. (1981). Autoshaping in the rat: Effects of omission on the form of the response. *Journal of the Experimental Analysis of Behavior* **36**: 75-91.
- Davis, M. (1992). The role of the amygdala in conditioned fear. In *The amygdala: neurobiological aspects of emotion, memory, and mental dysfunction* (Aggleton, J. P., ed.), pp. 255-306. Wiley-Liss, New York.
- Davis, M. (1997). Neurobiology of fear responses: the role of the amygdala. *Journal of Neuropsychiatry and Clinical Neuroscience* **9**: 382-402.
- Davison, M. C. (1987). The analysis of concurrent-chain performance. In *Quantitative Analyses of Behavior: V. The Effect of Delay and of Intervening Events on Reinforcement Value* (Commons, M. L., Mazur, J. E., Nevin, J. A. & Rachlin, H., eds.), pp. 225-241. Lawrence Erlbaum, Hillsdale, New Jersey.
- Dawson, G. R. & Dickinson, A. (1990). Performance on ratio and interval schedules with matched reinforcement rates. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **42**: 225-239.
- de Borchgrave, R. (1995). On the involvement of the rat nucleus accumbens in instrumental performance. Unpublished PhD thesis, University of Oxford.
- de Villiers, A. S., Russell, V. A., Sagvolden, T., Searson, A., Jaffer, A. & Taljaard, J. J. (1995). Alpha 2-adrenoceptor mediated inhibition of [3H]dopamine release from nucleus accumbens slices and monoamine levels in a rat model for attention-deficit hyperactivity disorder. *Neurochemical Research* **20**: 427-433.
- de Villiers, P. A. & Herrnstein, R. J. (1976). Toward a law of response strength. *Psychological Bulletin* **83**: 1131-1153.
- Dehaene, S., Posner, M. I. & Tucker, D. M. (1994). Commentary: Localization of a neural system for error detection and compensation. *Psychological Science* **5**: 303-305.
- DeLong, M. R. & Georgopoulos, A. P. (1981). Motor functions of the basal ganglia. In *Handbook of Physiology, Section 1, The Nervous System, Vol. 2, Part 2* (Brookhart, J. M., Mountcastle, V. B. & Brooks, V. B., eds.), pp. 1017-1061. American Physiological Society, Bethesda, Maryland.

- Deniau, J. M., Kitai, S. T., Donoghue, J. P. & Grofova, I. (1982). Neuronal interactions in the substantia nigra pars reticulata through axon collaterals of the projection neurons. An electrophysiological and morphological study. *Experimental Brain Research* **47**: 105-113.
- Derbyshire, S. W., Vogt, B. A. & Jones, A. K. (1998). Pain and Stroop interference tasks activate separate processing modules in anterior cingulate cortex. *Experimental Brain Research* **118**: 52-60.
- Devinsky, O., Morrell, M. J. & Vogt, B. A. (1995). Contributions of anterior cingulate cortex to behaviour. *Brain* **118**: 279-306.
- Di Chiara, G. (1998). A motivational learning hypothesis of the role of mesolimbic dopamine in compulsive drug use. *Journal of Psychopharmacology* **12**: 54-67.
- Dickinson, A. (1980). *Contemporary animal learning theory*, Cambridge University Press, Cambridge.
- Dickinson, A. (1985). Actions and habits - the development of behavioral autonomy. *Philosophical Transactions of the Royal Society of London, Series B - Biological Sciences* **308**: 67-78.
- Dickinson, A. (1986). Re-examination of the role of the instrumental contingency in the sodium-appetite irrelevant incentive effect. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **38**: 161-172.
- Dickinson, A. (1994). Instrumental conditioning. In *Animal Learning and Cognition* (Mackintosh, N. J., ed.), pp. 45-79. Academic Press, San Diego.
- Dickinson, A. & Balleine, B. (1994). Motivational control of goal-directed action. *Animal Learning & Behavior* **22**: 1-18.
- Dickinson, A., Balleine, B., Watt, A., Gonzalez, F. & Boakes, R. A. (1995). Motivational control after extended instrumental training. *Animal Learning & Behavior* **23**: 197-206.
- Dickinson, A. & Dawson, G. R. (1987a). The role of the instrumental contingency in the motivational control of performance. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **39**: 77-93.
- Dickinson, A. & Dawson, G. R. (1987b). Pavlovian processes in the motivational control of instrumental performance. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **39**: 201-213.
- Dickinson, A. & Dearing, M. F. (1979). Appetitive-aversive interactions and inhibitory processes. In *Mechanisms of learning and motivation* (Dickinson, A. & Boakes, R. A., eds.), pp. 203-231. Erlbaum, Hillsdale, New Jersey.
- Dickinson, A., Nicholas, D. J. & Adams, C. D. (1983). The effect of the instrumental training contingency on susceptibility to reinforcer devaluation. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **35**: 35-51.
- Dickinson, A., Smith, J. & Mirenowicz, J. (2000). Dissociation of Pavlovian and instrumental incentive learning under dopamine antagonists. *Behavioral Neuroscience* **114**: 468-483.
- Dickinson, A., Watt, A. & Griffiths, W. J. H. (1992). Free-operant acquisition with delayed reinforcement. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **45**: 241-258.
- Dietrich, A. & Allen, J. D. (1998). Functional dissociation of the prefrontal cortex and the hippocampus in timing behavior. *Behavioral Neuroscience* **112**: 1043-1047.
- Dix, S. L., Coutureau, E. & Killcross, A. S. (2000). The role of the nucleus accumbens in fear-related behaviours in rats. *European Journal of Neuroscience* **12** (supplement 11): 157.
- Doty, R. W. & Bosma, J. F. (1956). An electromyographic analysis of reflex deglutition. *Journal of Neurophysiology* **19**: 44-60.
- Drevets, W. C. (2000). PET imaging correlates of depression and mania. *Journal of Psychopharmacology* **14** (supplement): A1.
- Drevets, W. C., Price, J. L., Simpson, J. R., Todd, R. D., Reich, T., Vannier, M. & Raichle, M. E. (1997). Subgenual prefrontal cortex abnormalities in mood disorders. *Nature* **386**: 824-827.
- Dum, R. P. & Strick, P. L. (1993). Cingulate motor areas. In *Neurobiology of cingulate cortex and limbic thalamus: a comprehensive handbook* (Vogt, B. A. & Gabriel, M., eds.), pp. 415-441. Birkhäuser, Boston, Massachusetts.
- Dunn, R., Williams, B. & Royalty, P. (1987). Devaluation of stimuli contingent on choice: evidence for conditioned reinforcement. *Journal of the Experimental Analysis of Behavior* **48**: 117-131.
- Dunnett, S. B. & Iversen, S. D. (1982). Neurotoxic lesions of ventrolateral but not anteromedial neostriatum in rats impair differential reinforcement of low rates (DRL) performance. *Behavioural Brain Research* **6**: 213-226.

- Ernst, M., Zametkin, A. J., Matochik, J. A., Jons, P. H. & Cohen, R. M. (1998). DOPA decarboxylase activity in attention deficit hyperactivity disorder adults. A [fluorine-18]fluorodopa positron emission tomographic study. *Journal of Neuroscience* **18**: 5901-5907.
- Estes, W. K. (1943). Discriminative conditioning. I. A discriminative property of conditioned anticipation. *Journal of Experimental Psychology* **32**: 150-155.
- Estes, W. K. (1948). Discriminative conditioning. II. Effects of a Pavlovian conditioned stimulus upon a subsequently established operant response. *Journal of Experimental Psychology* **38**: 173-177.
- Evenden, J. L. (1998). Serotonergic and steroidal influences on impulsive behaviour in rats. *Comprehensive Summaries of Uppsala Dissertations from the Faculty of Medicine* **764**.
- Evenden, J. L. (1999a). Impulsivity: a discussion of clinical and experimental findings. *Journal of Psychopharmacology* **13**: 180-192.
- Evenden, J. L. (1999b). Varieties of impulsivity. *Psychopharmacology* **146**: 348-361.
- Evenden, J. L. & Carli, M. (1985). The effects of 6-hydroxydopamine lesions of the nucleus accumbens and caudate nucleus of rats on feeding in a novel environment. *Behavioural Brain Research* **15**: 63-70.
- Evenden, J. L. & Meyerson, B. (1999). The behavior of spontaneously hypertensive and Wistar Kyoto rats under a paced fixed consecutive number schedule of reinforcement. *Pharmacology, Biochemistry and Behavior* **63**: 71-82.
- Evenden, J. L. & Ryan, C. N. (1996). The pharmacology of impulsive behaviour in rats: the effects of drugs on response choice with varying delays of reinforcement. *Psychopharmacology* **128**: 161-170.
- Evenden, J. L. & Ryan, C. N. (1999). The pharmacology of impulsive behaviour in rats VI: the effects of ethanol and selective serotonergic drugs on response choice with varying delays of reinforcement. *Psychopharmacology* **146**: 413-421.
- Everitt, B. J. (1990). Sexual motivation: a neural and behavioural analysis of the mechanisms underlying appetitive and copulatory responses of male rats. *Neuroscience and Biobehavioral Reviews* **14**: 217-232.
- Everitt, B. J., Cador, M. & Robbins, T. W. (1989). Interactions between the amygdala and ventral striatum in stimulus-reward associations: studies using a second-order schedule of sexual reinforcement. *Neuroscience* **30**: 63-75.
- Everitt, B. J., Cardinal, R. N., Hall, J., Parkinson, J. A. & Robbins, T. W. (2000a). Differential involvement of amygdala subsystems in appetitive conditioning and drug addiction. In *The amygdala: a functional analysis*, Second edition (Aggleton, J. P., ed.), pp. 353-390. Oxford University Press, New York.
- Everitt, B. J., Fray, P., Kostarczyk, E., Taylor, S. & Stacey, P. (1987). Studies of instrumental behavior with sexual reinforcement in male rats (*Rattus norvegicus*): I. Control by brief visual stimuli paired with a receptive female. *Journal of Comparative Psychology* **101**: 395-406.
- Everitt, B. J., Morris, K. A., O'Brien, A. & Robbins, T. W. (1991). The basolateral amygdala-ventral striatal system and conditioned place preference: further evidence of limbic-striatal interactions underlying reward-related processes. *Neuroscience* **42**: 1-18.
- Everitt, B. J., Parkinson, J. A., Lachenal, G., Halkerston, K. M., Rudarakanchana, N., Cardinal, R. N., Hall, J., Morrison, C. H., Dalley, J. W., Howes, S. R. & Robbins, T. W. (2000b). Effects of limbic corticostriatal lesions on autoshaping performance in rats. *Society for Neuroscience Abstracts* **26**: 979.
- Everitt, B. J., Parkinson, J. A., Olmstead, M. C., Arroyo, M., Robledo, P. & Robbins, T. W. (1999). Associative processes in addiction and reward - The role of amygdala-ventral striatal subsystems. *Annals of the New York Academy of Sciences* **877**: 412-438.
- Everitt, B. J. & Robbins, T. W. (1992). Amygdala-ventral striatal interactions and reward-related processes. In *The amygdala: neurobiological aspects of emotion, memory, and mental dysfunction* (Aggleton, J. P., ed.), pp. 401-430. Wiley-Liss, New York.
- Everitt, B. J. & Stacey, P. (1987). Studies of instrumental behavior with sexual reinforcement in male rats (*Rattus norvegicus*): II. Effects of preoptic area lesions, castration, and testosterone. *Journal of Comparative Psychology* **101**: 407-419.
- Eysenck, S. G. B. (1993). The I7: development of a measure of impulsivity and its relationship to the superfactors of personality. In *The impulsive client: theory, research and treatment* (McCown, W. G., Johnson, J. L. & Shure, M. B., eds.). American Psychological Association, Washington DC.
- Eysenck, S. G. B. & Eysenck, H. J. (1977). The place of impulsiveness in a dimensional system of personality description. *British Journal of Social and Clinical Psychology* **16**: 57-68.

- Falkenstein, M., Hohnsbein, J., Hoormann, J. & Blanke, L. (1990). Effects of errors in choice reaction tasks on the ERP under focused and divided attention. In *Psychophysiological brain research* (Brunia, C. H. M., Gaillard, A. W. K. & Kok, A., eds.), pp. 192-195. Tilburg University Press, Tilburg, The Netherlands.
- Falkenstein, M., Hoormann, J., Christ, S. & Hohnsbein, J. (2000). ERP components on reaction errors and their functional significance: a tutorial. *Biological Psychology* **51**: 87-107.
- Fallon, J. H. & Loughlin, S. E. (1995). Substantia nigra. In *The Rat Nervous System* (Paxinos, G., ed.), pp. 215-237. Academic Press, London.
- Feldman, R. S., Meyer, J. S. & Quenzer, L. F. (1997). *Principles of neuropsychopharmacology*, Sinauer, Sunderland, Massachusetts.
- Fisk, G. D. & Wyss, J. M. (1997). Pressor and depressor sites are intermingled in the cingulate cortex of the rat. *Brain Research* **754**: 204-212.
- Fletcher, P. J., Korth, K. M., Sabijan, M. S. & DeSousa, N. J. (1998). Injections of d-amphetamine into the ventral pallidum increase locomotor activity and responding for conditioned reward: a comparison with injections into the nucleus accumbens. *Brain Research* **805**: 29-40.
- Floresco, S. B., Yang, C. R., Phillips, A. G. & Blaha, C. D. (1998). Basolateral amygdala stimulation evokes glutamate receptor-dependent dopamine efflux in the nucleus accumbens of the anaesthetized rat. *European Journal of Neuroscience* **10**: 1241-1251.
- Freeman, J. H., Jr., Cuppernell, C., Flannery, K. & Gabriel, M. (1996). Limbic thalamic, cingulate cortical and hippocampal neuronal correlates of discriminative approach learning in rabbits. *Behavioural Brain Research* **80**: 123-136.
- Frith, C. D., Friston, K., Liddle, P. F. & Frackowiak, R. S. (1991). Willed action and the prefrontal cortex in man: a study with PET. *Proceedings of the Royal Society of London, Series B - Biological Sciences* **244**: 241-246.
- Fudge, J. L. & Haber, S. N. (2000). The central nucleus of the amygdala projection to dopamine subpopulations in primates. *Neuroscience* **97**: 479-494.
- Fuster, J. M. (1995). *Memory in the cerebral cortex: an empirical approach to neural networks in the human and nonhuman primate*, MIT Press, Cambridge, Massachusetts.
- Gabriel, M. (1990). Functions of anterior and posterior cingulate cortex during avoidance learning in rabbits. *Progress in Brain Research* **85**: 467-482.
- Gabriel, M. (1993). Discriminative avoidance learning: a model system. In *Neurobiology of cingulate cortex and limbic thalamus: a comprehensive handbook* (Vogt, B. A. & Gabriel, M., eds.), pp. 478-523. Birkhauser, Boston, Massachusetts.
- Gabriel, M., Foster, K., Orona, E., Saltwick, S. & Stanton, M. (1980a). Neuronal activity of cingulate cortex, anteroventral thalamus, and hippocampal formation in discriminative conditioning: encoding and extraction of the significance of conditional stimuli. In *Progress in Psychobiology and Physiological Psychology* (Sprague, J. M. & Epstein, A. N., eds.), Vol. 9, pp. 125-231. Academic Press, New York.
- Gabriel, M., Kang, E., Poremba, A., Kubota, Y., Allen, M. T., Miller, D. P. & Steinmetz, J. E. (1996). Neural substrates of discriminative avoidance learning and classical eyeblink conditioning in rabbits: a double dissociation. *Behavioural Brain Research* **82**: 23-30.
- Gabriel, M., Kubota, Y., Sparenborg, S., Straube, K. & Vogt, B. A. (1991a). Effects of cingulate cortical lesions on avoidance learning and training-induced unit activity in rabbits. *Experimental Brain Research* **86**: 585-600.
- Gabriel, M., Lambert, R. W., Foster, K., Orona, E., Sparenborg, S. & Maiorca, R. R. (1983). Anterior thalamic lesions and neuronal activity in the cingulate and retrosplenial cortices during discriminative avoidance behavior in rabbits. *Behavioral Neuroscience* **97**: 675-696.
- Gabriel, M., Miller, J. D. & Saltwick, S. E. (1977). Unit activity in cingulate cortex and anteroventral thalamus of the rabbit during differential conditioning and reversal. *Journal of Comparative and Physiological Psychology* **91**: 423-433.
- Gabriel, M. & Orona, E. (1982). Parallel and serial processes of the prefrontal and cingulate cortical systems during behavioral learning. *Brain Research Bulletin* **8**: 781-785.
- Gabriel, M., Orona, E., Foster, K. & Lambert, R. W. (1980b). Cingulate cortical and anterior thalamic neuronal correlates of reversal learning in rabbits. *Journal of Comparative and Physiological Psychology* **94**: 1087-1100.
- Gabriel, M., Sparenborg, S. P. & Stolar, N. (1987). Hippocampal control of cingulate cortical and anterior thalamic information processing during learning in rabbits. *Experimental Brain Research* **67**: 131-152.
- Gabriel, M., Vogt, B. A., Kubota, Y., Poremba, A. & Kang, E. (1991b). Training-stage related neuronal plasticity in limbic thalamus and cingulate cortex during learning: a possible key to mnemonic retrieval. *Behavioural Brain Research* **46**: 175-185.

- Gaffan, D. (1992). Amnesia for complex naturalistic scenes and for objects following fornix transection in the rhesus monkey. *European Journal of Neuroscience* **4**: 381-388.
- Gaffan, D. & Harrison, S. (1989). Place memory and scene memory: effects of fornix transection in the monkey. *Experimental Brain Research* **74**: 202-212.
- Gallagher, M. & Holland, P. C. (1992). Understanding the function of the central nucleus: is simple conditioning enough? In *The amygdala: neurobiological aspects of emotion, memory, and mental dysfunction* (Aggleton, J. P., ed.), pp. 307-321. Wiley-Liss, New York.
- Gallagher, M. & Holland, P. C. (1994). The amygdala complex: multiple roles in associative learning and attention. *Proceedings of the National Academy of Sciences of the United States of America* **91**: 11771-11776.
- Gallagher, M., McMahan, R. W. & Schoenbaum, G. (1999). Orbitofrontal cortex and representation of incentive value in associative learning. *Journal of Neuroscience* **19**: 6610-6614.
- Gallistel, C. R. (1994). Space and time. In *Animal Learning and Cognition* (Mackintosh, N. J., ed.), pp. 221-253. Academic Press, San Diego.
- Garavan, H., Pankiewicz, J., Bloom, A., Cho, J. K., Sperry, L., Ross, T. J., Salmeron, B. J., Risinger, R., Kelley, D. & Stein, E. A. (2000). Cue-induced cocaine craving: Neuroanatomical specificity for drug users and drug stimuli. *American Journal of Psychiatry* **157**: 1789-1798.
- Garcia, J. (1989). Food for Tolman: Cognition and cathexis in concert. In *Aversion, avoidance and anxiety* (Archer, T. & Nilsson, L.-G., eds.), pp. 45-85. Erlbaum, Hillsdale, New Jersey.
- Garrud, P., Goodall, G. & Mackintosh, N. J. (1981). Overshadowing of a stimulus-reinforcer association by an instrumental response. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **33**: 123-135.
- Gawin, F. H. (1991). Cocaine addiction: psychology and neurophysiology. *Science* **251**: 1580-1586.
- Gehring, W. J., Coles, M. G. H., Meyer, D. E. & Donchin, E. (1990). The error-related negativity: an event-related brain potential accompanying errors. *Psychophysiology* **27** (supplement): S34.
- Gehring, W. J., Goss, B., Coles, M. G. H., Meyer, D. E. & Donchin, E. (1993). A neural system for error detection and compensation. *Psychological Science* **4**: 385-390.
- Gehring, W. J., Himle, J. & Nisenson, L. G. (2000). Action monitoring dysfunction in obsessive-compulsive disorder. *Psychological Science* **11**: 1-6.
- Gehring, W. J. & Knight, R. T. (2000). Prefrontal-cingulate interactions in action monitoring. *Nature Neuroscience* **3**: 516-520.
- Gelder, M., Gath, D., Mayou, R. & Cowen, P. (1995). *Oxford Textbook of Psychiatry*. Third edition, Oxford University Press, Oxford.
- Gemba, H., Sasaki, K. & Brooks, V. B. (1986). 'Error' potentials in limbic cortex (anterior cingulate area 24) of monkeys during motor learning. *Neuroscience Letters* **70**: 223-227.
- Gentry, G. D., Weiss, B. & Laties, V. G. (1983). The microanalysis of fixed-interval responding. *Journal of the Experimental Analysis of Behavior* **39**: 327-343.
- Gerfen, C. R. (1992a). The neostriatal mosaic: multiple levels of compartmental organization. *Trends in Neurosciences* **15**: 133-139.
- Gerfen, C. R. (1992b). The neostriatal mosaic: multiple levels of compartmental organization in the basal ganglia. *Annual Review of Neuroscience* **15**: 285-320.
- Gewirtz, J. C. & Davis, M. (1998). Application of Pavlovian higher-order conditioning to the analysis of the neural substrates of fear conditioning. *Neuropharmacology* **37**: 453-459.
- Gibbon & Church (1981). Time left: linear versus logarithmic subjective time. *Journal of Experimental Psychology: Animal Behavior Processes* **7**: 87-108.
- Gibbon, J., Malapani, C., Dale, C. L. & Gallistel, C. (1997). Toward a neurobiology of temporal cognition: advances and challenges. *Current Opinion in Neurobiology* **7**: 170-184.
- Gödel, K. (1931). Über formal unentscheidbare Sätze der Principia mathematica und verwandter Systeme I. *Monatshefte für Mathematik und Physik* **38**: 173-198.
- Goodale, M. A. & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences* **15**: 20-25.
- Gottman, J. M. (1981). *Time-series analysis: a comprehensive introduction for social scientists*, Cambridge University Press, Cambridge.
- Grace, A. A. (1987). An electrophysiologist's eye view of the basal ganglia. *Behavioral and Brain Sciences* **10**: 214-215.

- Grace, R. C. (1996). Choice between fixed and variable delays to reinforcement in the adjusting-delay procedure and concurrent chains. *Journal of Experimental Psychology: Animal Behavior Processes* **22**: 362-383.
- Grice, G. R. (1948). The relation of secondary reinforcement to delayed reward in visual discrimination learning. *Journal of Experimental Psychology* **38**: 1-16.
- Grill, H. J. & Berridge, K. C. (1985). Taste reactivity as a measure of the neural control of palatability. *Progress in Psychobiology and Physiological Psychology* **11**: 1-61.
- Grindley, G. C. (1932). The formation of a simple habit in guinea pigs. *British Journal of Psychology* **23**: 127-147.
- Grossman, M., Crino, P., Reivich, M., Stern, M. B. & Hurtig, H. I. (1992). Attention and sentence processing deficits in Parkinson's disease: the role of anterior cingulate cortex. *Cerebral Cortex* **2**: 513-525.
- Grove, E. A., Domesick, V. B. & Nauta, W. J. (1986). Light microscopic evidence of striatal input to intrapallidal neurons of cholinergic cell group Ch4 in the rat: a study employing the anterograde tracer Phaseolus vulgaris leucoagglutinin (PHA-L). *Brain Research* **367**: 379-384.
- Guthrie, E. R. (1935). *The psychology of learning*, Harper, New York.
- Haber, S. N., Fudge, J. L. & McFarland, N. R. (2000). Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. *Journal of Neuroscience* **20**: 2369-2382.
- Hall, F. S., Wilkinson, L. S., Humby, T., Inglis, W., Kendall, D. A., Marsden, C. A. & Robbins, T. W. (1998). Isolation rearing in rats: pre- and postsynaptic changes in striatal dopaminergic systems. *Pharmacology, Biochemistry and Behavior* **59**: 859-872.
- Hall, J. (1999). The roles of the amygdala and hippocampus in Pavlovian conditioning. Unpublished PhD thesis, University of Cambridge.
- Hall, J., Parkinson, J. A., Connor, T. M. F., Di Ciano, P., Dickinson, A. & Everitt, B. J. (1999). The role of amygdala-ventral striatal sub-systems in Pavlovian to instrumental transfer. *Society for Neuroscience Abstracts* **25**: 90.
- Halliday, G., Harding, A. & Paxinos, G. (1995). Serotonin and tachykinin systems. In *The Rat Nervous System* (Paxinos, G., ed.), pp. 929-974. Academic Press, London.
- Han, J. S., Holland, P. C. & Gallagher, M. (1999). Disconnection of the amygdala central nucleus and substantia innominata/nucleus basalis disrupts increments in conditioned stimulus processing in rats. *Behavioral Neuroscience* **113**: 143-151.
- Han, J. S., McMahan, R. W., Holland, P. & Gallagher, M. (1997). The role of an amygdalo-nigrostriatal pathway in associative learning. *Journal of Neuroscience* **17**: 3913-3919.
- Hare, S. R. (1996). Low frequency climate variability and salmon production. Unpublished PhD thesis, University of Washington.
- Harlow, J. M. (1868). Recovery from the passage of an iron bar through the head. *Publications of the Massachusetts Medical Society* **2**: 327-347.
- Harmer, C. J. & Phillips, G. D. (1999). Enhanced dopamine efflux in the amygdala by a predictive, but not a non-predictive, stimulus: Facilitation by prior repeated D-amphetamine. *Neuroscience* **90**: 119-130.
- Harrison, A. A., Everitt, B. J. & Robbins, T. W. (1999). Central serotonin depletion impairs both the acquisition and performance of a symmetrically reinforced go/no-go conditional visual discrimination. *Behavioural Brain Research* **100**: 99-112.
- Hart, M., Poremba, A. & Gabriel, M. (1997). The nomadic engram: overtraining eliminates the impairment of discriminative avoidance behavior produced by limbic thalamic lesions. *Behavioural Brain Research* **82**: 169-177.
- Hatfield, T., Han, J. S., Conley, M., Gallagher, M. & Holland, P. (1996). Neurotoxic lesions of basolateral, but not central, amygdala interfere with Pavlovian second-order conditioning and reinforcer devaluation effects. *Journal of Neuroscience* **16**: 5256-5265.
- Hauber, W., Bohn, I. & Gierler, C. (2000). NMDA, but not dopamine D(2), receptors in the rat nucleus accumbens are involved in guidance of instrumental behavior by stimuli predicting reward magnitude. *Journal of Neuroscience* **20**: 6282-6288.
- Hearst, E. & Jenkins, H. M. (1974). *Sign tracking: the stimulus-reinforcer relation and directed action*, The Psychonomic Society, Austin, Texas.
- Heckerman, D. E., Horvitz, E. J. & Nathwani, B. N. (1992). Toward normative expert systems: Part I. The Pathfinder project. *Methods of Information in Medicine* **31**: 90-105.
- Heidbreder, C. A., Weiss, I. C., Domeney, A. M., Pryce, C., Homberg, J., Hedou, G., Feldon, J., Moran, M. C. & Nelson, P. (2000). Behavioral, neurochemical and endocrinological characterization of the early social isolation syndrome. *Neuroscience* **100**: 749-768.
- Heimer, L., Zahm, D. S. & Alheid, G. F. (1995). Basal ganglia. In *The Rat Nervous System* (Paxinos, G., ed.), pp. 579-628. Academic Press, London.

- Henderson, R. W., Patterson, J. M. & Jackson, R. L. (1980). Acquisition and retention of control of instrumental behavior by a cue signaling airblast: how specific are conditioned anticipations? *Learning & Motivation* **11**: 407-426.
- Herrnstein, R. J. (1961). Relative and absolute strength of responses as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior* **4**: 267-272.
- Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior* **13**: 243-266.
- Hershberger, W. A. (1986). An approach through the looking glass. *Animal Learning & Behavior* **14**: 443-451.
- Hetherington, A. W. & Ranson, S. W. (1939). Experimental hypothalamohypophyseal obesity in the rat. *Proceedings of the Society for Experimental Biology and Medicine* **41**: 465-466.
- Heyman, G. M. (1996). Resolving the contradictions of addiction. *Behavioral and Brain Sciences* **19**: 561-610.
- Hill, R. T. (1970). Facilitation of conditioned reinforcement as a mechanism of psychomotor stimulation. In *International Symposium on Amphetamines and Related Compounds* (Costa, E. & Garattini, S., eds.), pp. 781-795. Raven Press, New York.
- Ho, M., Mobini, S., Chiang, T., Bradshaw, C. M. & Szabadi, E. (1999). Theory and method in the quantitative analysis of "impulsive choice" behaviour: implications for psychopharmacology. *Psychopharmacology* **146**: 362-372.
- Ho, M. Y., Al-Zahrani, S. S., Velazquez Martinez, D. N., Lopez Cabrera, M., Bradshaw, C. M. & Szabadi, E. (1995). The role of the ascending 5-hydroxytryptaminergic pathways in timing behaviour: further observations with the interval bisection task. *Psychopharmacology* **120**: 213-219.
- Ho, M. Y., Wogar, M. A., Bradshaw, C. M. & Szabadi, E. (1997). Choice between delayed reinforcers: Interaction between delay and deprivation level. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **50**: 193-202.
- Holland, P. C. (1981). Acquisition of representation-mediated conditioned food aversions. *Learning & Motivation* **12**: 1-18.
- Holland, P. C. (1990a). Event representation in Pavlovian conditioning: image and action. *Cognition* **37**: 105-131.
- Holland, P. C. (1990b). Forms of memory in Pavlovian conditioning. In *Brain organization and memory: cells, systems and circuits* (McGaugh, J. L., Weinberger, N. M. & Lynch, N. M., eds.), pp. 78-105. Oxford University Press, New York.
- Holland, P. C. (1997). Brain mechanisms for changes in processing of conditioned stimuli in Pavlovian conditioning: Implications for behavior theory. *Animal Learning & Behavior* **25**: 373-399.
- Holland, P. C. (1998). Amount of training affects associatively-activated event representation. *Neuropharmacology* **37**: 461-469.
- Holland, P. C. & Gallagher, M. (1993a). Amygdala central nucleus lesions disrupt increments, but not decrements, in conditioned stimulus processing. *Behavioral Neuroscience* **107**: 246-253.
- Holland, P. C. & Gallagher, M. (1993b). Effects of amygdala central nucleus lesions on blocking and unblocking. *Behavioral Neuroscience* **107**: 235-245.
- Holland, P. C. & Gallagher, M. (1999). Amygdala circuitry in attentional and representational processes. *Trends in Cognitive Sciences* **3**: 65-73.
- Holland, P. C., Han, J. S. & Gallagher, M. (2000). Lesions of the amygdala central nucleus alter performance on a selective attention task. *Journal of Neuroscience* **20**: 6701-6706.
- Holman, J. G. & Mackintosh, N. J. (1981). The control of appetitive instrumental responding does not depend on classical conditioning to the discriminative stimulus. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **33**: 21-31.
- Holroyd, C., Reichler, J. & Coles, M. G. H. (1999). Is the error-related negativity generated by a dopaminergic error signal for reinforcement learning? Hypothesis and model. *Journal of Cognitive Neuroscience* **11** (supplement): 45.
- Holroyd, C. B., Dien, J. & Coles, M. G. (1998). Error-related scalp potentials elicited by hand and foot movements: evidence for an output-independent error-processing system in humans. *Neuroscience Letters* **242**: 65-68.
- Holyoak, K. J. & Spellman, B. A. (1993). Thinking. *Annual Review of Psychology* **44**: 265-315.
- Houk, J. C., Adams, J. L. & Barto, A. G. (1995). A model of how the basal ganglia generate and use neural signals that predict reinforcement. In *Models of information processing in the basal ganglia* (Houk, J. C., Davis, J. L. & Beiser, D. G., eds.), pp. 249-270. MIT Press, Cambridge, Massachusetts.
- Houk, J. C. & Wise, S. P. (1995). Distributed modular architectures linking basal ganglia, cerebellum, and cerebral cortex: their role in planning and controlling action. *Cerebral Cortex* **2**: 95-110.

- Howell, D. C. (1997). *Statistical Methods for Psychology*. Fourth edition, Wadsworth, Belmont, California.
- Howes, S. R., Dalley, J. W., Morrison, C. H., Robbins, T. W. & Everitt, B. J. (2000). Leftward shift in the acquisition of cocaine self-administration in isolation-reared rats: relationship to extracellular levels of dopamine, serotonin and glutamate in the nucleus accumbens and amygdala-striatal FOS expression. *Psychopharmacology* **151**: 55-63.
- Hsu, M. M. & Shyu, B. C. (1997). Electrophysiological study of the connection between medial thalamus and anterior cingulate cortex in the rat. *Neuroreport* **8**: 2701-2707.
- Hull, C. L. (1932). The goal gradient hypothesis and maze learning. *Psychological Review* **39**: 25-43.
- Hull, C. L. (1943). *Principles of behavior*, Appleton-Century-Crofts, New York.
- Huynh, H. & Feldt, L. S. (1970). Conditions under which mean square ratios in repeated measures designs have exact *F*-distributions. *Journal of the American Statistical Association* **65**: 1582-1589.
- Ito, M. (1985). Choice and amount of reinforcement in rats. *Learning & Motivation* **16**: 95-108.
- Ito, R., Dalley, J. W., Howes, S. R., Robbins, T. W. & Everitt, B. J. (2000). Dissociation in conditioned dopamine release in the nucleus accumbens core and shell in response to cocaine cues and during cocaine-seeking behavior in rats. *Journal of Neuroscience* **20**: 7489-7495.
- Jaspers, R., Schwarz, M., Sontag, K. H. & Cools, A. R. (1984). Caudate nucleus and programming behaviour in cats: role of dopamine in switching motor patterns. *Behavioral Brain Research* **14**: 17-28.
- Jenkins, H. M. & Moore, B. R. (1973). The form of the auto-shaped response with food or water reinforcers. *Journal of the Experimental Analysis of Behavior* **20**: 163-181.
- Jiménez-Castellanos, J. & Graybiel, A. M. (1989). Evidence that histochemically distinct zones of the primate substantia nigra pars compacta are related to patterned distributions of nigrostriatal projection neurons and striatonigral fibres. *Experimental Brain Research* **74**: 227-238.
- Jones, G. H., Hernandez, T. D., Kendall, D. A., Marsden, C. A. & Robbins, T. W. (1992). Dopaminergic and serotonergic function following isolation rearing in rats: study of behavioural responses and postmortem and in vivo neurochemistry. *Pharmacology, Biochemistry and Behavior* **43**: 17-35.
- Kacelnik, A. (1997). Normative and descriptive models of decision making: time discounting and risk sensitivity. In *Characterizing human psychological adaptations (Ciba Foundation Symposium 208)*, pp. 51-70. Wiley, Chichester.
- Kahneman, D., Slovic, P. & Tversky, A., Eds. (1982). *Judgement Under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press.
- Kalivas, P. W., Pierce, R. C., Cornish, J. & Sorg, B. A. (1998). A role for sensitization in craving and relapse in cocaine addiction. *Journal of Psychopharmacology* **12**: 49-53.
- Kamin, L. J. (1968). 'Attention-like' processes in classical conditioning. In *Miami Symposium on the Prediction of Behavior: Aversive Stimulation* (Jones, M. R., ed.), pp. 9-33. University of Miami Press, Miami.
- Kamin, L. K. (1969). Predictability, surprise, attention and conditioning. In *Punishment and Aversive Behavior* (Campbell, B. A. & Church, R. M., eds.), pp. 279-296. Appleton-Century-Crofts, New York.
- Kelley, A. E., SmithRoe, S. L. & Holahan, M. R. (1997). Response-reinforcement learning is dependent on N-methyl-D-aspartate receptor activation in the nucleus accumbens core. *Proceedings of the National Academy of Sciences of the United States of America* **94**: 12174-12179.
- Kelley, A. E. & Stinus, L. (1985). Disappearance of hoarding behavior after 6-hydroxydopamine lesions of the mesolimbic dopamine neurons and its reinstatement with l-dopa. *Behavioral Neuroscience* **99**: 531-545.
- Kemp, J. M. & Powell, T. P. S. (1971). The connections of the striatum and globus pallidus: synthesis and speculation. *Philosophical Transactions of the Royal Society of London, Series B - Biological Sciences* **262**: 441-457.
- Keppel, G. (1982). *Design and analysis: a researcher's handbook*. Second edition, Englewood Cliffs: Prentice-Hall, London.
- Keppel, G. (1991). *Design and analysis: a researcher's handbook*. Third edition, Prentice-Hall, London.
- Killcross, A. S., Everitt, B. J. & Robbins, T. W. (1997a). Symmetrical effects of amphetamine and alpha-flupenthixol on conditioned punishment and conditioned reinforcement: contrasts with midazolam. *Psychopharmacology* **129**: 141-152.
- Killcross, A. S., Everitt, B. J. & Robbins, T. W. (1998). Dissociable effects of excitotoxic lesions of amygdala sub-nuclei on appetitive conditioning. *Journal of Psychopharmacology* **12 (supplement A)**: A4.

- Killcross, S., Robbins, T. W. & Everitt, B. J. (1997b). Different types of fear-conditioned behaviour mediated by separate nuclei within amygdala. *Nature* **388**: 377-380.
- Killeen, P. R. (1972). The matching law. *Journal of the Experimental Analysis of Behavior* **17**: 489-495.
- Killeen, P. R. (1981). Averaging theory. In *Quantification of steady-state operant behaviour* (Bradshaw, C. M. & Szabadi, E., eds.), pp. 21-34. Elsevier, Amsterdam.
- Killeen, P. R. & Fetterman, J. G. (1988). A behavioral theory of timing. *Psychological Review* **95**: 274-295.
- Kish, G. B. (1966). Studies of sensory reinforcement. In *Operant behavior: Areas of research and application* (Honig, W. K., ed.), pp. 109-159. Appleton-Century-Crofts, New York.
- Konorski, J. (1948). *Conditioned reflexes and neuron organization*, Cambridge University Press, Cambridge.
- Konorski, J. (1967). *Integrative activity of the brain*, University of Chicago Press, Chicago.
- Koob, G. F. (1992). Dopamine, addiction and reward. *Seminars in the Neurosciences* **4**: 139-148.
- Koob, G. F., Riley, S. J., Smith, S. C. & Robbins, T. W. (1978). Effects of 6-hydroxydopamine lesions of the nucleus accumbens septi and olfactory tubercle on feeding, locomotor activity, and amphetamine anorexia in the rat. *Journal of Comparative and Physiological Psychology* **92**: 917-927.
- Kopp, B. & Wolff, M. (2000). Brain mechanisms of selective learning: event-related potentials provide evidence for error-driven learning in humans. *Biological Psychology* **51**: 223-246.
- Krettek, J. E. & Price, J. L. (1977). The cortical projections of the mediodorsal nucleus and adjacent thalamic nuclei in the rat. *Journal of Comparative Neurology* **171**: 157-192.
- Krieg, W. J. S. (1946). Connections of the cerebral cortex. I. The albino rat. A topography of the cortical areas. *Journal of Comparative Neurology* **84**: 221-275.
- Kupferman, I. (1991). Hypothalamus and limbic system: motivation. In *Principles of Neural Science*, Third edition (Kandel, E. R., Schwartz, J. H. & Jessell, T. M., eds.), pp. 750-760. Appleton-Lange, East Norwalk, Connecticut.
- Laties, V. G. (1972). The modification of drug effects on behaviour by external discriminative stimuli. *Journal of Pharmacology and Experimental Therapeutics* **183**: 1-13.
- Laties, V. G. & Weiss, B. (1966). Influence of drugs on behavior controlled by internal and external stimuli. *Journal of Pharmacology and Experimental Therapeutics* **152**: 388-396.
- Laties, V. G., Wood, R. W. & Rees, D. C. (1981). Stimulus control and the effects of d-amphetamine in the rat. *Psychopharmacology* **75**: 277-282.
- Lattal, K. A. (1987). Considerations in the experimental analysis of reinforcement delay. In *Quantitative Analyses of Behavior: V. The Effect of Delay and of Intervening Events on Reinforcement Value* (Commons, M. L., Mazur, J. E., Nevin, J. A. & Rachlin, H., eds.), pp. 107-123. Lawrence Erlbaum, Hillsdale, New Jersey.
- Lattal, K. A. & Gleeson, S. (1990). Response acquisition with delayed reinforcement. *Journal of Experimental Psychology: Animal Behavior Processes* **16**: 27-39.
- LeDoux, J. E. (1992). Emotion and the amygdala. In *The amygdala: neurobiological aspects of emotion, memory, and mental dysfunction* (Aggleton, J. P., ed.), pp. 339-352. Wiley-Liss, New York.
- LeDoux, J. E. (2000). Emotion circuits in the brain. *Annual Review of Neuroscience* **23**: 155-184.
- Linnoila, M., Virkkunen, M., George, T. & Higley, D. (1993). Impulse control disorders. *International Clinical Psychopharmacology* **8 (Supplement 1)**: 53-56.
- Linnoila, M., Virkkunen, M., Scheinin, M., Nuutila, A., Rimon, R. & Goodwin, F. K. (1983). Low cerebrospinal fluid 5-hydroxyindoleacetic acid concentration differentiates impulsive from nonimpulsive violent behavior. *Life Sciences* **33**: 2609-2614.
- Liotti, M., Woldorff, M. G., Perez, R. & Mayberg, H. S. (2000). An ERP study of the temporal course of the Stroop color-word interference effect. *Neuropsychologia* **38**: 701-711.
- Lorenz (1939). Vergleichende Verhaltensforschung [Comparative behaviour research]. *Zoologischer Anzeiger Supplement* **12**: 69-109.

- Lovibond, P. F. (1983). Facilitation of instrumental behavior by a Pavlovian appetitive conditioned stimulus. *Journal of Experimental Psychology: Animal Behavior Processes* **9**: 225-247.
- Luu, P., Collins, P. & Tucker, D. M. (2000a). Mood, personality, and self-monitoring: negative affect and emotionality in relation to frontal lobe mechanisms of error monitoring. *Journal of Experimental Psychology: General* **129**: 43-60.
- Luu, P., Flaisch, T. & Tucker, D. M. (2000b). Medial frontal cortex in action monitoring. *Journal of Neuroscience* **20**: 464-469.
- Lyon, M. & Robbins, T. W. (1975). The action of central nervous system stimulant drugs: a general theory concerning amphetamine effects. In *Current Developments in Psychopharmacology* (Essman, W. B. & Valzelli, L., eds.), Vol. 2, pp. 79-163. Spectrum Publications, New York.
- Maas, L. C., Lukas, S. E., Kaufman, M. J., Weiss, R. D., Daniels, S. L., Rogers, V. W., Kukes, T. J. & Renshaw, P. F. (1998). Functional magnetic resonance imaging of human brain activation during cue-induced cocaine craving. *American Journal of Psychiatry* **155**: 124-126.
- MacDonald, A. W., Cohen, J. D., Stenger, V. A. & Carter, C. S. (2000). Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. *Science* **288**: 1835-1838.
- Mackintosh, N. J. (1974). *The Psychology of Animal Learning*, Academic Press, London.
- Mackintosh, N. J. (1983). *Conditioning and associative learning*, Oxford University Press, Oxford.
- MacLean, P. D. (1949). Psychosomatic disease and the "visceral brain": recent developments bearing on the Papez theory of emotion. *Psychosomatic Medicine* **11**: 338-353.
- MacLean, P. D. (1952). Some psychiatric implications of physiological studies on the frontotemporal portion of the limbic system (visceral brain). *Electroencephalography and Clinical Neurophysiology* **4**: 407-418.
- MacLean, P. D. (1993). Introduction: perspectives on cingulate cortex in the limbic system. In *Neurobiology of cingulate cortex and limbic thalamus: a comprehensive handbook* (Vogt, B. A. & Gabriel, M., eds.), pp. 1-15. Birkhauser, Boston, Massachusetts.
- MacLeod, C. M. & MacDonald, P. A. (2000). Interdimensional interference in the Stroop effect: uncovering the cognitive and neural anatomy of attention. *Trends in Cognitive Sciences* **4**: 383-391.
- Maldonado-Irizarry, C. S. & Kelley, A. E. (1995). Excitotoxic lesions of the core and shell subregions of the nucleus accumbens differentially disrupt body-weight regulation and motor activity in the rat. *Brain Research Bulletin* **38**: 551-559.
- Málková, L., Gaffan, D. & Murray, E. A. (1997). Excitotoxic lesions of the amygdala fail to produce impairment in visual learning for auditory secondary reinforcement but interfere with reinforcer devaluation effects in rhesus monkeys. *Journal of Neuroscience* **17**: 6011-6020.
- Maren, S. (1999). Neurotoxic or electrolytic lesions of the ventral subiculum produce deficits in the acquisition and expression of Pavlovian fear conditioning in rats. *Behavioral Neuroscience* **113**: 283-290.
- Mark, T. A. & Gallistel, C. R. (1994). Kinetics of matching. *Journal of Experimental Psychology: Animal Behavior Processes* **20**: 79-95.
- Marsden, C. D. (1992). Dopamine and basal ganglia disorders in humans. *Seminars in the Neurosciences* **4**: 171-178.
- Marshall, J. F. & Teitelbaum, P. (1977). New considerations in the neuropsychology of motivated behaviors. In *Handbook of Psychopharmacology* (Iversen, L. L., Iversen, S. D. & Snyder, S. H., eds.), Vol. 7, pp. 201-222. Plenum Press, New York.
- Mayberg, H. S. (1997). Limbic-cortical dysregulation: A proposed model of depression. *Journal of Neuropsychiatry and Clinical Neurosciences* **9**: 471-481.
- Mayberg, H. S. (2000). The neuropsychopathology of mood regulation and depression: converging evidence from PET studies of unipolar and bipolar patients. *Journal of Psychopharmacology* **14** (supplement): A1.
- Mayberg, H. S., Brannan, S. K., Mahurin, R. K., Brickman, J. S., Jerabek, P. A., Martin, C. C. & Fox, P. T. (1996). Anterior cingulate function and mood: Evidence from FDG PET studies of primary and secondary depression. *Neurology* **46**: 28004-28004.
- Mayberg, H. S., Brannan, S. K., Mahurin, R. K., Jerabek, P. A., Brickman, J. S., Tekell, J. L., Silva, J. A., McGinnis, S., Glass, T. G., Martin, C. C. & Fox, P. T. (1997). Cingulate function in depression: A potential predictor of treatment response. *Neuroreport* **8**: 1057-1061.
- Mayberg, H. S., Lewis, P. J., Regenold, W. & Wagner, H. N. (1994). Paralimbic hypoperfusion in unipolar depression. *Journal of Nuclear Medicine* **35**: 929-934.
- Mayberg, H. S., Liotti, M., Brannan, S. K., McGinnis, S., Mahurin, R. K., Jerabek, P. A., Silva, J. A., Tekell, J. L., Martin, C. C., Lancaster, J. L. & Fox, P. T. (1999). Reciprocal limbic-cortical function and negative mood: Converging PET findings in depression and normal sadness. *American Journal of Psychiatry* **156**: 675-682.

- Mazur, J. E. (1984). Tests of an equivalence rule for fixed and variable reinforcer delays. *Journal of Experimental Psychology: Animal Behavior Processes* **10**: 426-436.
- Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In *Quantitative Analyses of Behavior: V. The Effect of Delay and of Intervening Events on Reinforcement Value* (Commons, M. L., Mazur, J. E., Nevin, J. A. & Rachlin, H., eds.), pp. 55-73. Lawrence Erlbaum, Hillsdale, New Jersey.
- Mazur, J. E. (1988). Estimation of indifference points with an adjusting-delay procedure. *Journal of the Experimental Analysis of Behavior* **49**: 37-47.
- Mazur, J. E. (1991). Choice with probabilistic reinforcement - effects of delay and conditioned reinforcers. *Journal of the Experimental Analysis of Behavior* **55**: 63-77.
- Mazur, J. E. (1995). Conditioned reinforcement and choice with delayed and uncertain primary reinforcers. *Journal of the Experimental Analysis of Behavior* **63**: 139-150.
- Mazur, J. E. (1997). Choice, delay, probability, and conditioned reinforcement. *Animal Learning & Behavior* **25**: 131-147.
- Mazur, J. E. (2000). Tradeoffs among delay, rate, and amount of reinforcement. *Behavioural Processes* **49**: 1-10.
- Mazur, J. E., Stellar, J. R. & Waraczynski, M. (1987). Self-control choice with electrical stimulation of the brain as a reinforcer. *Behavioural Processes* **15**: 143-153.
- McAlonan, G. M., Robbins, T. W. & Everitt, B. J. (1993). Effects of medial dorsal thalamic and ventral pallidal lesions on the acquisition of a conditioned place preference: further evidence for the involvement of the ventral striatopallidal system in reward-related processes. *Neuroscience* **52**: 605-620.
- McCleary, R. & Hay, R. A. (1980). *Applied time series analysis for the social sciences*, Sage, London.
- McCullough, L. D., Cousins, M. S. & Salamone, J. D. (1993). The role of nucleus accumbens dopamine in responding on a continuous reinforcement operant schedule: a neurochemical and behavioral study. *Pharmacology, Biochemistry and Behavior* **46**: 581-586.
- McDonald, A. J. (1998). Cortical pathways to the mammalian amygdala. *Progress in Neurobiology* **55**: 257-332.
- McGeorge, A. J. & Faull, R. L. M. (1989). The organization of the projection from the cerebral cortex to the striatum in the rat. *Neuroscience* **29**: 503-537.
- Mead, A. N. & Stephens, D. N. (1998). AMPA-receptors are involved in the expression of amphetamine-induced behavioural sensitisation, but not in the expression of amphetamine-induced conditioned activity in mice. *Neuropharmacology* **37**: 1131-1138.
- Meck, W. H. (1983). Selective adjustment of the speed of internal clock and memory processes. *Journal of Experimental Psychology: Animal Behavior Processes* **9**: 171-201.
- Mehlman, P. T., Higley, J. D., Faucher, I., Lilly, A. A., Taub, D. M., Vickers, J., Suomi, S. J. & Linnoila, M. (1994). Low CSF 5-HIAA concentrations and severe aggression and impaired impulse control in nonhuman primates. *American Journal of Psychiatry* **151**: 1485-1491.
- Mérö, L. (1998). *Moral Calculations: Game Theory, Logic, and Human Frailty*, Springer-Verlag, New York.
- Meunier, M. & Destrade, C. (1986). [Paradoxical transitory facilitation of performance in the Hebb-Williams labyrinth after lesion of the cingulate cortex in mice]. *Comptes Rendus de l'Academie des Sciences. Serie III, Sciences de la Vie* **302**: 43-46.
- Meunier, M. & Destrade, C. (1988). Electrolytic but not ibotenic acid lesions of the posterior cingulate cortex produce transitory facilitation of learning in mice. *Behavioural Brain Research* **27**: 161-172.
- Meunier, M., Jaffard, R. & Destrade, C. (1991). Differential involvement of anterior and posterior cingulate cortices in spatial discriminative learning in a T-maze in mice. *Behavioural Brain Research* **44**: 133-143.
- Miltner, W. H. R., Braun, C. H. & Coles, M. G. H. (1997). Event-related brain potentials following incorrect feedback in a time-estimation task: Evidence for a "generic" neural system for error detection. *Journal of Cognitive Neuroscience* **9**: 788-798.
- Mirenovic, J. & Schultz, W. (1994). Importance of unpredictability for reward responses in primate dopamine neurons. *Journal of Neurophysiology* **72**: 1024-1027.
- Mirenovic, J. & Schultz, W. (1996). Preferential activation of midbrain dopamine neurons by appetitive rather than aversive stimuli. *Nature* **379**: 449-451.

- Mishkin, M., Malamut, B. & Bachevalier, J. (1984). Memories and habits: two neural systems. In *Neurobiology of Learning and Memory* (Lynch, G., McGaugh, J. L. & Weinberger, N. M., eds.), pp. 65-77. Guildford Press, New York.
- Mitchell, S. H. (1999). Measures of impulsivity in cigarette smokers and non-smokers. *Psychopharmacology* **146**: 455-464.
- Mittleman, G., Brener, J. & Robbins, T. W. (1990). Physiological correlates of schedule-induced activities in rats. *American Journal of Physiology* **259**: R485-491.
- Moerschbaecher, J. M., Boren, J. J., Schrot, J. & Fontes, J. C. (1979). Effects of cocaine and d-amphetamine on the repeated acquisition and performance of conditional discriminations. *Journal of the Experimental Analysis of Behavior* **31**: 127-140.
- Mogenson, G. J., Jones, D. L. & C.Y., Y. (1980). From motivation to action: functional interface between the limbic system and the motor system. *Progress in Neurobiology* **14**: 69-97.
- Montague, P. R., Dayan, P. & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *Journal of Neuroscience* **16**: 1936-1947.
- Morgan, M. A. & Ledoux, J. E. (1995). Differential contribution of dorsal and ventral medial prefrontal cortex to the acquisition and extinction of conditioned fear in rats. *Behavioral Neuroscience* **109**: 681-688.
- Morgan, M. J. & Nicholas, D. J. (1979). Discrimination between reinforced action patterns in the rat. *Learning & Motivation* **10**: 1-22.
- Morrissey, G., Ho, M. Y., Wogar, M. A., Bradshaw, C. M. & Szabadi, E. (1994). Effect of lesions of the ascending 5-hydroxytryptaminergic pathways on timing behaviour investigated with the fixed-interval peak procedure. *Psychopharmacology* **114**: 463-468.
- Morrissey, G., Wogar, M. A., Bradshaw, C. M. & Szabadi, E. (1993). Effect of lesions of the ascending 5-hydroxytryptaminergic pathways on timing behaviour investigated with an interval bisection task. *Psychopharmacology* **112**: 80-85.
- Mowrer, O. H. (1960). *Learning Theory and Behavior*, Wiley, New York.
- Muir, J. L., Everitt, B. J. & Robbins, T. W. (1996). The cerebral cortex of the rat and visual attentional function: dissociable effects of medio-frontal, cingulate, anterior dorsolateral, and parietal cortex lesions on a five-choice serial reaction time task. *Cerebral Cortex* **6**: 470-481.
- Mullen, R. J., Buck, C. R. & Smith, A. M. (1992). NeuN, a neuronal specific nuclear protein in vertebrates. *Development* **116**: 201-211.
- Murray, E. A. & Bussey, T. J. (1999). Perceptual-mnemonic functions of the perirhinal cortex. *Trends in Cognitive Sciences* **3**: 142-151.
- Murray, E. A., Bussey, T. J., Hampton, R. R. & Saksida, L. M. (2000). The parahippocampal region and object identification. *Annals of the New York Academy of Sciences* **911**: 166-174.
- Myerson, J. & Green, L. (1995). Discounting of delayed rewards: models of individual choice. *Journal of the Experimental Analysis of Behavior* **64**: 263-276.
- Nauta, W. J. H. (1971). The problem of the frontal lobe: a reinterpretation. *Journal of Psychiatric Research* **8**: 167-187.
- Neafsey, E. J., Terreberry, K. M., Hurley, K. M., Ruit, K. G. & Fryszak, R. J. (1993). Anterior cingulate cortex in rodents: connections, visceral control functions, and implications for emotion. In *Neurobiology of cingulate cortex and limbic thalamus: a comprehensive handbook* (Vogt, B. A. & Gabriel, M., eds.), pp. 206-233. Birkhäuser, Boston, Massachusetts.
- Neave, N., Lloyd, S., Sahgal, A. & Aggleton, J. P. (1994). Lack of effect of lesions in the anterior cingulate cortex and retrosplenial cortex on certain tests of spatial memory in the rat. *Behavioural Brain Research* **65**: 89-101.
- Nicholson, D. A. & Freeman, J. H. (2000). Lesions of the perirhinal cortex impair sensory preconditioning in rats. *Behavioural Brain Research* **112**: 69-75.
- Norgren, R. (1995). Gustatory system. In *The Rat Nervous System* (Paxinos, G., ed.), pp. 751-771. Academic Press, London.
- Norman, D. A. & Shallice, T. (1986). Attention to action: willed and automatic control of behaviour. In *Consciousness and Self-Regulation (Volume 4): Advances in Research and Theory* (Davidson, R. J., Schwartz, G. E. & Shapiro, D., eds.), pp. 1-18. Plenum Press, New York.
- Oakley, D. A. & Russell, I. S. (1972). Neocortical lesions and Pavlovian conditioning. *Physiology and Behavior* **8**: 915-922.
- Oakley, D. A. & Russell, I. S. (1975). Role of cortex on Pavlovian discrimination learning. *Physiology and Behavior* **15**: 315-321.
- Oakley, D. A. & Russell, I. S. (1976). Subcortical nature of Pavlovian differentiation in the rabbit. *Physiology and Behavior* **17**: 947-954.
- O'Brien, C. P., Childress, A. R., Ehrman, R. & Robbins, S. J. (1998). Conditioning factors in drug abuse: can they explain compulsion? *Journal of Psychopharmacology* **12**: 15-22.

- O'Brien, C. P., Ehrman, R. N. & Ternes, J. W. (1986). Classical conditioning in human opioid dependence. In *Behavioural analysis of drug dependence* (Goldberg, S. R. & Stolerman, I. P., eds.), pp. 329-356. Academic Press, London.
- Öngür, D., An, X. & Price, J. L. (1998). Prefrontal cortical projections to the hypothalamus in macaque monkeys. *Journal of Comparative Neurology* **401**: 480-505.
- Öngür, D. & Price, J. L. (2000). The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. *Cerebral Cortex* **10**: 206-219.
- Packard, M. G., Hirsh, R. & White, N. M. (1989). Differential effects of fornix and caudate nucleus lesions on two radial maze tasks: evidence for multiple memory systems. *Journal of Neuroscience* **9**: 1465-1472.
- Packard, M. G. & McGaugh, J. L. (1996). Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiology of Learning and Memory* **65**: 65-72.
- Packard, M. G. & White, N. M. (1991). Dissociation of hippocampus and caudate nucleus memory systems by posttraining intracerebral injection of dopamine agonists. *Behavioral Neuroscience* **105**: 295-306.
- Papa, M., Sagvolden, T., Sergeant, J. A. & Sadile, A. G. (1996). Reduced CaMKII-positive neurones in the accumbens shell of an animal model of attention-deficit hyperactivity disorder. *Neuroreport* **7**: 3017-3020.
- Papa, M., Sellitti, S. & Sadile, A. G. (2000). Remodeling of neural networks in the anterior forebrain of an animal model of hyperactivity and attention deficits as monitored by molecular imaging probes. *Neuroscience and Biobehavioral Reviews* **24**: 149-156.
- Papa, M., Sergeant, J. A. & Sadile, A. G. (1998). Reduced transduction mechanisms in the anterior accumbal interface of an animal model of attention-deficit hyperactivity disorder. *Behavioural Brain Research* **94**: 187-195.
- Papez, J. W. (1937). A proposed mechanism of emotion. *Archives of Neurology and Psychiatry* **38**: 725-743.
- Pardo, J. V., Pardo, P. J., Janer, K. W. & Raichle, M. E. (1990). The anterior cingulate cortex mediates processing selection in the Stroop attentional conflict paradigm. *Proceedings of the National Academy of Sciences of the United States of America* **87**: 256-259.
- Parkinson, J. A. (1998). Limbic corticostriatal circuitry underlying Pavlovian associative learning. Unpublished PhD thesis, University of Cambridge.
- Parkinson, J. A., Buckby, T. B., Zandi, M. S., Robbins, T. W. & Everitt, B. J. (1999a). Nucleus accumbens lesions enhance the acquisition of a simple instrumental visual discrimination. *Society for Neuroscience Abstracts* **25**: 90.
- Parkinson, J. A., Cardinal, R. N. & Everitt, B. J. (2000a). Limbic cortical-ventral striatal systems underlying appetitive conditioning. *Progress in Brain Research* **126**: 263-285.
- Parkinson, J. A., Dalley, J. W., Cardinal, R. N., Bamford, A., Fehner, B., Lachenal, G., Rudarakanchana, N., Halkerston, K. M., Robbins, T. W. & Everitt, B. J. (submitted). Nucleus accumbens dopamine depletion impairs both acquisition and performance of appetitive Pavlovian approach behaviour: implications for mesoaccumbens dopamine function. .
- Parkinson, J. A., Olmstead, M. C., Burns, L. H., Robbins, T. W. & Everitt, B. J. (1999b). Dissociation in effects of lesions of the nucleus accumbens core and shell on appetitive Pavlovian approach behavior and the potentiation of conditioned reinforcement and locomotor activity by d-amphetamine. *Journal of Neuroscience* **19**: 2401-2411.
- Parkinson, J. A., Robbins, T. W. & Everitt, B. J. (1996). Lesions of the nucleus accumbens core, but not basolateral amygdala or subiculum, disrupt stimulus-reward learning in a novel autoshaping procedure. *Society for Neuroscience Abstracts* **22**: 1118.
- Parkinson, J. A., Robbins, T. W. & Everitt, B. J. (1999c). Selective excitotoxic lesions of the nucleus accumbens core and shell differentially affect aversive Pavlovian conditioning to discrete and contextual cues. *Psychobiology* **27**: 256-266.
- Parkinson, J. A., Robbins, T. W. & Everitt, B. J. (2000b). Dissociable roles of the central and basolateral amygdala in appetitive emotional learning. *European Journal of Neuroscience* **12**: 405-413.
- Parkinson, J. A., Willoughby, P. J., Robbins, T. W. & Everitt, B. J. (2000c). Disconnection of the anterior cingulate cortex and nucleus accumbens core impairs Pavlovian approach behavior: Further evidence for limbic cortical-ventral striatopallidal systems. *Behavioral Neuroscience* **114**: 42-63.
- Paus, T., Petrides, M., Evans, A. C. & Meyer, E. (1993). Role of the human anterior cingulate cortex in the control of oculomotor, manual, and speech responses: a positron emission tomography study. *Journal of Neurophysiology* **70**: 453-469.

- Pavlov, I. P. (1927). *Conditioned Reflexes*, Oxford University Press, Oxford.
- Paxinos, G. & Watson, C. (1996). *The Rat Brain in Stereotaxic Coordinates*. Compact Third (CD-Rom) edition, Academic Press.
- Paxinos, G. & Watson, C. (1998). *The Rat Brain in Stereotaxic Coordinates*. Fourth edition, Academic Press.
- Pearce, J. M. & Hall, G. (1980). A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review* **106**: 532-552.
- Pennartz, C. M., Ameerun, R. F., Groenewegen, H. J. & Lopes da Silva, F. H. (1993). Synaptic plasticity in an in vitro slice preparation of the rat nucleus accumbens. *European Journal of Neuroscience* **5**: 107-117.
- Pennartz, C. M., Groenewegen, H. J. & Lopes da Silva, F. H. (1994). The nucleus accumbens as a complex of functionally distinct neuronal ensembles: an integration of behavioural, electrophysiological and anatomical data. *Progress in Neurobiology* **42**: 719-761.
- Pennartz, C. M. & Kitai, S. T. (1991). Hippocampal inputs to identified neurons in an in vitro slice preparation of the rat nucleus accumbens: evidence for feed-forward inhibition. *Journal of Neuroscience* **11**: 2838-2847.
- Pennartz, C. M. A. (1995). The ascending neuromodulatory systems in learning by reinforcement: Comparing computational conjectures with experimental findings. *Brain Research Reviews* **21**: 219-245.
- Petersen, S. E., van Mier, H., Fiez, J. A. & Raichle, M. E. (1998). The effects of practice on the functional anatomy of task performance. *Proceedings of the National Academy of Sciences of the United States of America* **95**: 853-860.
- Phillips, R. G. & LeDoux, J. E. (1992). Differential contribution of amygdala and hippocampus to cued and contextual fear conditioning. *Behavioral Neuroscience* **106**: 274-285.
- Pinker, S. (1997). *How The Mind Works*, Allen Lane - The Penguin Press, London.
- Pitkänen, A. (2000). Connectivity of the rat amygdaloid complex. In *The amygdala: a functional analysis*, Second edition (Aggleton, J. P., ed.), pp. 31-115. Oxford University Press, New York.
- Poremba, A. & Gabriel, M. (1997). Amygdalar lesions block discriminative avoidance learning and cingulothalamic training-induced neuronal plasticity in rabbits. *Journal of Neuroscience* **17**: 5237-5244.
- Poremba, A. & Gabriel, M. (1999). Amygdala neurons mediate acquisition but not maintenance of instrumental avoidance behavior in rabbits. *Journal of Neuroscience* **19**: 9635-9641.
- Posner, M. I. (1995). Attention in cognitive neuroscience: an overview. In *The Cognitive Neurosciences* (Gazzaniga, M. S., ed.), pp. 615-624. MIT Press, Cambridge.
- Poulos, C. X., Le, A. D. & Parker, J. L. (1995). Impulsivity predicts individual susceptibility to high levels of alcohol self-administration. *Behavioural Pharmacology* **6**: 810-814.
- Powell, D. A., Watson, K. & Maxwell, B. (1994). Involvement of subdivisions of the medial prefrontal cortex in learned cardiac adjustments in rabbits. *Behavioral Neuroscience* **108**: 294-307.
- Price, J. L. (1995). Thalamus. In *The Rat Nervous System* (Paxinos, G., ed.), pp. 629-648. Academic Press, London.
- Rachlin, H. (1971). On the tautology of the matching law. *Journal of the Experimental Analysis of Behavior* **15**: 249-251.
- Rachlin, H., Raineri, A. & Cross, D. (1991). Subjective probability and delay. *Journal of the Experimental Analysis of Behavior* **55**: 233-244.
- Ragozzino, M. E., Adams, S. & Kesner, R. P. (1998). Differential involvement of the dorsal anterior cingulate and prelimbic-infralimbic areas of the rodent prefrontal cortex in spatial working memory. *Behavioral Neuroscience* **112**: 293-303.
- Ragozzino, M. E., Wilcox, C., Raso, M. & Kesner, R. P. (1999). Involvement of rodent prefrontal cortex subregions in strategy switching. *Behavioral Neuroscience* **113**: 32-41.
- Raichle, M. E., Fiez, J. A., Videen, T. O., MacLeod, A. M., Pardo, J. V., Fox, P. T. & Petersen, S. E. (1994). Practice-related changes in human brain functional anatomy during nonmotor learning. *Cerebral Cortex* **4**: 8-26.
- Reading, P. J. & Dunnett, S. B. (1991). The effects of excitotoxic lesions of the nucleus accumbens on a matching to position task. *Behavioural Brain Research* **46**: 17-29.
- Reading, P. J., Dunnett, S. B. & Robbins, T. W. (1991). Dissociable roles of the ventral, medial and lateral striatum on the acquisition and performance of a complex visual stimulus-response habit. *Behavioural Brain Research* **45**: 147-161.

- Redgrave, P., Prescott, T. J. & Gurney, K. (1999a). The basal ganglia: A vertebrate solution to the selection problem? *Neuroscience* **89**: 1009-1023.
- Redgrave, P., Prescott, T. J. & Gurney, K. (1999b). Is the short-latency dopamine response too short to signal reward error? *Trends in Neurosciences* **22**: 146-151.
- Rescorla, R. A. (1990a). Evidence for an association between the discriminative stimulus and the response-outcome association in instrumental learning. *Journal of Experimental Psychology: Animal Behavior Processes* **16**: 326-334.
- Rescorla, R. A. (1990b). Instrumental responses become associated with reinforcers that differ in one feature. *Animal Learning & Behavior* **18**: 206-211.
- Rescorla, R. A. (1990c). The role of information about the response-outcome relation in instrumental discrimination learning. *Journal of Experimental Psychology: Animal Behavior Processes* **16**: 262-270.
- Rescorla, R. A. & Wagner, A. R. (1972). A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and non-reinforcement. In *Classical Conditioning II: Current Research and Theory* (Black, A. & Prokasy, W., eds.), pp. 64-99. Appleton-Century-Crofts, New York.
- Revusky, S. & Garcia, J. (1970). Learned associations over long delays. In *The Psychology of Learning and Motivation* (Bower, G. H., ed.), Vol. 4, pp. 1-84. Academic Press, New York.
- Richards, J. B., Chock, M. A., Carlson, B., de Wit, H. & Seiden, L. (1997a). Comparison of two models of impulsive behavior in rats: effects of amphetamine and haloperidol. *Society for Neuroscience Abstracts* **23**: 2406.
- Richards, J. B., Mitchell, S. H., de Wit, H. & Seiden, L. S. (1997b). Determination of discount functions in rats with an adjusting-amount procedure. *Journal of the Experimental Analysis of Behavior* **67**: 353-366.
- Richards, J. B., Sabol, K. E. & de Wit, H. (1999). Effects of methamphetamine on the adjusting amount procedure, a model of impulsive behavior in rats. *Psychopharmacology* **146**: 432-439.
- Richards, J. B. & Seiden, L. S. (1995). Serotonin depletion increases impulsive behavior in rats. *Society for Neuroscience Abstracts* **21**: 1693.
- Richardson, N. R. & Gratton, A. (1998). Changes in medial prefrontal cortical dopamine levels associated with response-contingent food reward: an electrochemical study in rat. *Journal of Neuroscience* **18**: 9130-9138.
- Robbins, T. W. (1976). Relationship between reward-enhancing and stereotypical effects of psychomotor stimulant drugs. *Nature* **264**: 57-59.
- Robbins, T. W. (1978). The acquisition of responding with conditioned reinforcement: effects of pipradrol, methylphenidate, d-amphetamine, and nomifensine. *Psychopharmacology* **58**: 79-87.
- Robbins, T. W. & Everitt, B. J. (1992). Functions of dopamine in the dorsal and ventral striatum. *Seminars in the Neurosciences* **4**: 119-127.
- Robbins, T. W. & Everitt, B. J. (1996). Neurobehavioural mechanisms of reward and motivation. *Current Opinion in Neurobiology* **6**: 228-236.
- Robbins, T. W., Giardini, V., Jones, G. H., Reading, P. & Sahakian, B. J. (1990a). Effects of dopamine depletion from the caudate-putamen and nucleus accumbens septi on the acquisition and performance of a conditional discrimination task. *Behavioural Brain Research* **38**: 243-261.
- Robbins, T. W., Hutcherson, D. M., Parkinson, J. A. & Everitt, B. J. (2000). The effects of nucleus accumbens core and shell lesions on intravenous heroin self-administration and the acquisition of heroin-seeking behaviour (under a second-order schedule). *European Journal of Neuroscience* **12** (supplement 11): 156.
- Robbins, T. W. & Koob, G. F. (1978). Pipradrol enhances reinforcing properties of stimuli paired with brain stimulation. *Pharmacology, Biochemistry and Behavior* **8**: 219-222.
- Robbins, T. W. & Koob, G. F. (1980). Selective disruption of displacement behaviour by lesions of the mesolimbic dopamine system. *Nature* **285**: 409-412.
- Robbins, T. W., Mittleman, G., O'Brien, J. & Winn, P. (1990b). The neuropsychological significance of stereotypy induced by stimulant drugs. In *Neurobiology of stereotyped behaviour* (Cooper, S. J. & Dourish, C. T., eds.), pp. 25-63. Oxford University Press, Oxford.
- Robbins, T. W. & Sahakian, B. J. (1983). Behavioural effects of psychomotor stimulant drugs: clinical and neuropsychological implications. In *Stimulants: Neurochemical, Behavioural, and Clinical Perspectives* (Creese, I., ed.), pp. 301-338. Raven Press, New York.
- Robbins, T. W., Watson, B. A., Gaskin, M. & Ennis, C. (1983). Contrasting interactions of pipradrol, d-amphetamine, cocaine, cocaine analogues, apomorphine and other drugs with conditioned reinforcement. *Psychopharmacology* **80**: 113-119.

- Roberts, A. C., Robbins, T. W. & Weiskrantz, L. (1998). Discussion and conclusions. In *The prefrontal cortex: executive and cognitive functions* (Roberts, A. C., Robbins, T. W. & Weiskrantz, L., eds.), pp. 221-242. Oxford University Press, Oxford.
- Roberts, S. (1981). Isolation of an internal clock. *Journal of Experimental Psychology: Animal Behavior Processes* **7**: 242-268.
- Robinson, T. E. & Berridge, K. C. (1993). The neural basis of drug craving: an incentive-sensitization theory of addiction. *Brain Research Reviews* **18**: 247-291.
- Robledo, P., Robbins, T. W. & Everitt, B. J. (1996). Effects of excitotoxic lesions of the central amygdaloid nucleus on the potentiation of reward-related stimuli by intra-accumbens amphetamine. *Behavioral Neuroscience* **110**: 981-990.
- Rogers, R. D., Owen, A. M., Middleton, H. C., Williams, E. J., Pickard, J. D., Sahakian, B. J. & Robbins, T. W. (1999). Choosing between small, likely rewards and large, unlikely rewards activates inferior and orbital prefrontal cortex. *Journal of Neuroscience* **19**: 9029-9038.
- Rolls, E. T. (1999). *The Brain and Emotion*, Oxford University Press, Oxford.
- Rolls, E. T. (2000). Orbitofrontal cortex and reward. *Cerebral Cortex* **10**: 284-294.
- Rubia, K., Overmeyer, S., Taylor, E., Brammer, M., Williams, S. C. R., Simmons, A. & Bullmore, E. T. (1999). Hypofrontality in attention deficit hyperactivity disorder during higher-order motor control: A study with functional MRI. *American Journal of Psychiatry* **156**: 891-896.
- Russell, S. J. & Norvig, P. N. (1995). *Artificial Intelligence: a modern approach*, Prentice-Hall, Upper Saddle River, New Jersey.
- Russell, V., de Villiers, A., Sagvolden, T., Lamm, M. & Taljaard, J. (1998). Differences between electrically-, ritalin- and D-amphetamine-stimulated release of [H-3]dopamine from brain slices suggest impaired vesicular storage of dopamine in an animal model of attention-deficit hyperactivity disorder. *Behavioural Brain Research* **94**: 163-171.
- Russell, V., Devilliers, A., Sagvolden, T., Lamm, M. & Taljaard, J. (1995). Altered dopaminergic function in the prefrontal cortex, nucleus accumbens and caudate-putamen of an animal model of attention-deficit hyperactivity disorder - the spontaneously hypertensive rat. *Brain Research* **676**: 343-351.
- Russell, V. A. (2000). The nucleus accumbens motor-limbic interface of the spontaneously hypertensive rat as studied in vitro by the superfusion slice technique. *Neuroscience and Biobehavioral Reviews* **24**: 133-136.
- Sadile, A. G. (2000). Multiple evidence of a segmental defect in the anterior forebrain of an animal model of hyperactivity and attention deficit. *Neuroscience and Biobehavioral Reviews* **24**: 161-169.
- Sagvolden, T. (2000). Behavioral validation of the spontaneously hypertensive rat (SHR) as an animal model of attention-deficit/hyperactivity disorder (AD/HD). *Neuroscience and Biobehavioral Reviews* **24**: 31-39.
- Sagvolden, T., Aase, H., Zeiner, P. & Berger, D. (1998). Altered reinforcement mechanisms in attention-deficit/hyperactivity disorder. *Behavioural Brain Research* **94**: 61-71.
- Sagvolden, T., Metzger, M. A., Schiorbeck, H. K., Rugland, A. L., Spinnangr, I. & Sagvolden, G. (1992). The spontaneously hypertensive rat (SHR) as an animal model of childhood hyperactivity (ADHD): changed reactivity to reinforcers and to psychomotor stimulants. *Behavioral and Neural Biology* **58**: 103-112.
- Sagvolden, T., Pettersen, M. B. & Larsen, M. C. (1993). Spontaneously hypertensive rats (SHR) as a putative animal model of childhood hyperkinesis: SHR behavior compared to four other rat strains. *Physiology and Behavior* **54**: 1047-1055.
- Sagvolden, T. & Sergeant, J. A. (1998). Attention deficit/hyperactivity disorder - from brain dysfunctions to behaviour. *Behavioural Brain Research* **94**: 1-10.
- Saksida, L. M. & Bussey, T. J. (1998). Toward a neural network model of visual object identification in primate inferotemporal cortex. *Society for Neuroscience Abstracts* **24**: 1906.
- Saksida, L. M., Bussey, T. J. & Murray, E. A. (2000). One-pair discrimination learning in rhesus monkeys with perirhinal cortex lesions. *Society for Neuroscience Abstracts* **26**: 546.
- Salamone, J. D. (1994). The involvement of nucleus accumbens dopamine in appetitive and aversive motivation. *Behavioural Brain Research* **61**: 117-133.
- Salamone, J. D., Cousins, M. S. & Bucher, S. (1994). Anhedonia or anergia? Effects of haloperidol and nucleus accumbens dopamine depletion on instrumental response selection in a T-maze cost/benefit procedure. *Behavioural Brain Research* **65**: 221-229.

- Salamone, J. D., Cousins, M. S. & Snyder, B. J. (1997). Behavioral functions of nucleus accumbens dopamine: empirical and conceptual problems with the anhedonia hypothesis. *Neuroscience and Biobehavioral Reviews* **21**: 341-359.
- Salamone, J. D., Kurth, P. A., McCullough, L. D., Sokolowski, J. D. & Cousins, M. S. (1993). The role of brain dopamine in response initiation: effects of haloperidol and regionally specific dopamine depletions on the local rate of instrumental responding. *Brain Research* **628**: 218-226.
- Salamone, J. D., Steinpreis, R. E., McCullough, L. D., Smith, P., Grebel, D. & Mahan, K. (1991). Haloperidol and nucleus accumbens dopamine depletion suppress lever pressing for food but increase free food consumption in a novel food choice procedure. *Psychopharmacology* **104**: 515-521.
- Sarter, M. & Markowitsch, H. (1985). Involvement of the amygdala in learning and memory: a critical review with emphasis on anatomical relations. *Behavioral Neuroscience* **99**: 342-380.
- Scheffers, M. K. & Coles, M. G. H. (2000). Performance monitoring in a confusing world: Error-related brain activity, judgments of response accuracy, and types of errors. *Journal of Experimental Psychology: Human Perception and Performance* **26**: 141-151.
- Scheffers, M. K., Coles, M. G. H., Bernstein, P., Gehring, W. J. & Donchin, E. (1996). Event-related brain potentials and error-related processing: An analysis of incorrect responses to go and no-go stimuli. *Psychophysiology* **33**: 42-53.
- Scheffers, M. K., Humphrey, D. G., Stanny, R. R., Kramer, A. F. & Coles, M. G. H. (1999). Error-related processing during a period of extended wakefulness. *Psychophysiology* **36**: 149-157.
- Schoenbaum, G., Chiba, A. A. & Gallagher, M. (1998). Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nature Neuroscience* **1**: 155-159.
- Schoenbaum, G., Chiba, A. A. & Gallagher, M. (1999). Neural encoding in orbitofrontal cortex and basolateral amygdala during olfactory discrimination learning. *Journal of Neuroscience* **19**: 1876-1884.
- Schultz, W. (1994). Behavior-related activity of primate dopamine neurons. *Revue Neurologique (Paris)* **150**: 634-639.
- Schultz, W., Apicella, P. & Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. *Journal of Neuroscience* **13**: 900-913.
- Schultz, W., Apicella, P., Romo, R. & Scarnati, E. (1995a). Context-dependent activity in primate striatum reflecting past and future behavioral events. In *Models of Information Processing in the Basal Ganglia* (Houk, J. C., Davis, J. L. & Beiser, D. G., eds.), pp. 11-27. MIT Press, Cambridge, Massachusetts / London.
- Schultz, W., Apicella, P., Scarnati, E. & Ljungberg, T. (1992). Neuronal activity in monkey ventral striatum related to the expectation of reward. *Journal of Neuroscience* **12**: 4595-4610.
- Schultz, W., Dayan, P. & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science* **275**: 1593-1599.
- Schultz, W. & Dickinson, A. (2000). Neuronal coding of prediction errors. *Annual Review of Neuroscience* **23**: 473-500.
- Schultz, W., Romo, R., Ljungberg, T., Mirenowicz, J., Hollerman, J. R. & Dickinson, A. (1995b). Reward-related signals carried by dopamine neurons. In *Models of Information Processing in the Basal Ganglia* (Houk, J. C., Davis, J. L. & Beiser, D. G., eds.), pp. 233-248. MIT Press, Cambridge, Massachusetts.
- Schultz, W., Tremblay, L. & Hollerman, J. R. (1998). Reward prediction in primate basal ganglia and frontal cortex. *Neuropharmacology* **37**: 421-429.
- Schultz, W., Tremblay, W. & Hollerman, J. R. (2000). Reward processing in primate orbitofrontal cortex and basal ganglia. *Cerebral Cortex* **10**: 272-283.
- Schwartzbaum, J. S. (1965). Discrimination behavior after amygdectomy in monkeys: visual and somesthetic learning and perceptual capacity. *Journal of Comparative and Physiological Psychology* **60**: 314-319.
- Seamans, J. K., Floresco, S. B. & Phillips, A. G. (1995). Functional differences between the prelimbic and anterior cingulate regions of the rat prefrontal cortex. *Behavioral Neuroscience* **109**: 1063-1073.
- Selden, N. R., Everitt, B. J., Jarrard, L. E. & Robbins, T. W. (1991). Complementary roles for the amygdala and hippocampus in aversive conditioning to explicit and contextual cues. *Neuroscience* **42**: 335-350.

- Setlow, B., Holland, P. C. & Gallagher, M. (2000). Involvement of a basolateral amygdala complex-nucleus accumbens system in appetitive Pavlovian second-order conditioning. *Society for Neuroscience Abstracts* **26**: 1504.
- Setlow, B. & McGaugh, J. L. (1998). Sulpiride infused into the nucleus accumbens posttraining impairs memory of spatial water maze training. *Behavioral Neuroscience* **112**: 603-610.
- Setlow, B. & McGaugh, J. L. (1999). Differential effects of immediate posttraining sulpiride microinfusions into the nucleus accumbens shell and core on Morris water maze retention. *Psychobiology* **27**: 248-255.
- Shastri, L. & Ajjanagadde, V. (1993). From simple associations to systematic reasoning: a connectionist representation of rules, variables and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences* **16**: 417-494.
- Sheffield, F. D. (1965). Relation between classical conditioning and instrumental learning. In *Classical conditioning: a symposium* (Prokasy, W. F., ed.), pp. 302-322. Appleton-Century-Crofts, New York.
- Shettleworth, S. J. (1975). Reinforcement and the organization of behavior in golden hamsters: Hunger, environment, and food reinforcement. *Journal of Experimental Psychology: Animal Behavior Processes* **1**: 56-87.
- Shimp, C. P. (1967). The reinforcement of short interresponse times. *Journal of the Experimental Analysis of Behavior* **10**: 425-434.
- Shimp, C. P. (1969). Optimal behaviour in free-operant experiments. *Journal of Experimental Psychology: General* **108**: 48-67.
- Shizgal, P. (1997). Neural basis of utility estimation. *Current Opinion in Neurobiology* **7**: 198-208.
- Simerly, R. B. (1995). Anatomical substrates of hypothalamic integration. In *The Rat Nervous System* (Paxinos, G., ed.), pp. 353-376. Academic Press, London.
- Smith, J. W. & Dickinson, A. (1998). The dopamine antagonist, pimozone, abolishes Pavlovian-instrumental transfer. *Journal of Psychopharmacology* **12** (supplement A): A6.
- Sokolowski, J. D. & Salamone, J. D. (1998). The role of accumbens dopamine in lever pressing and response allocation: effects of 6-OHDA injected into core and dorsomedial shell. *Pharmacology, Biochemistry and Behavior* **59**: 557-566.
- Solanto, M. V. (1998). Neuropsychopharmacological mechanisms of stimulant drug action in attention-deficit hyperactivity disorder: a review and integration. *Behavioural Brain Research* **94**: 127-152.
- Sonuga-Barke, E. J. S., Saxton, T. & Hall, M. (1998). The role of interval underestimation in hyperactive children's failure to suppress responses over time. *Behavioural Brain Research* **94**: 45-50.
- Soubrié, P. (1986). Reconciling the role of central serotonin neurons in human and animal behavior. *Behavioral and Brain Sciences* **9**: 319-335.
- Soubrié, P., Blas, C., Ferron, A. & Glowinski, J. (1983). Chlordiazepoxide reduces in vivo serotonin release in the basal ganglia of encephale isolé but not anesthetized cats: evidence for a dorsal raphé site of action. *Journal of Pharmacology and Experimental Therapeutics* **226**: 526-532.
- Sougné, J. (1998). Connectionism and the problem of multiple instantiation. *Trends in Cognitive Sciences* **2**: 183-189.
- Spence, K. W. (1956). *Behavior Theory and Conditioning*, Prentice-Hall, Englewood Cliffs, New Jersey.
- StatSoft (1999). *Electronic Statistics Textbook* (URL: <<http://www.statsoft.com/textbook/stathome.html>>), StatSoft, Tulsa, Oklahoma.
- Steinmetz, J. E. (2000). Brain substrates of classical eyeblink conditioning: a highly localized but also distributed system. *Behavioural Brain Research* **110**: 13-24.
- Steinmetz, J. E., Sears, L. L., Gabriel, M., Kubota, Y. & Poremba, A. (1991). Cerebellar interpositus nucleus lesions disrupt classical nictitating-membrane conditioning but not discriminative avoidance learning in rabbits. *Behavioural Brain Research* **45**: 71-80.
- Stephens, D. N. (1995). A glutamatergic hypothesis of drug dependence - extrapolations from benzodiazepine receptor ligands. *Behavioural Pharmacology* **6**: 425-446.
- Stern, C. E. & Passingham, R. E. (1995). The nucleus accumbens in monkeys (*Macaca fascicularis*): III. Reversal learning. *Experimental Brain Research* **106**: 239-247.
- Stricker, E. M. & Zigmond, M. J. (1976). Recovery of function after damage to central catecholamine-containing neurons: a neurochemical model for the lateral hypothalamic syndrome. In *Progress in Psychology and Physiological Psychology* (Sprague, J. M. & Epstein, A. N., eds.), pp. 121-188. Academic Press, New York.
- Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology* **18**: 643-662.

- Stroustrup, B. (1986). *The C++ Programming Language*, Addison-Wesley, Reading, Massachusetts.
- Stroustrup, B. (2000). *The C++ Programming Language: Special Edition*, Addison-Wesley, Reading, Massachusetts.
- Sutton, R. S. (1988). Learning to predict by the method of temporal differences. *Machine Learning* **3**: 9-44.
- Swanson, L. W. (1987). The hypothalamus. In *Handbook of Chemical Neuroanatomy* (Björklund, A., Hökfelt, T. & Swanson, L. W., eds.), Vol. 5, pp. 1-124. Elsevier, Amsterdam.
- Szostak, C. & Tombaugh, T. N. (1981). Use of a fixed consecutive number schedule of reinforcement to investigate the effects of pimozide on behavior controlled by internal and external stimuli. *Pharmacology, Biochemistry and Behavior* **15**: 609-617.
- Tarpy, R. M. (1997). *Contemporary Learning Theory and Research*, McGraw-Hill, New York.
- Taylor, J. R. & Horger, B. A. (1999). Enhanced responding for conditioned reward produced by intra-accumbens amphetamine is potentiated after cocaine sensitization. *Psychopharmacology* **142**: 31-40.
- Taylor, J. R. & Robbins, T. W. (1984). Enhanced behavioural control by conditioned reinforcers following microinjections of d-amphetamine into the nucleus accumbens. *Psychopharmacology* **84**: 405-412.
- Taylor, J. R. & Robbins, T. W. (1986). 6-Hydroxydopamine lesions of the nucleus accumbens, but not of the caudate nucleus, attenuate enhanced responding with reward-related stimuli produced by intra-accumbens d-amphetamine. *Psychopharmacology* **90**: 390-397.
- Terrace, H. S. (1963). Errorless discrimination learning in the pigeon: effects of chlorpromazine and imipramine. *Science* **140**: 318-319.
- Thiébot, M. H. (1986). Are serotonergic neurons involved in the control of anxiety and in the anxiolytic activity of benzodiazepines? *Pharmacology, Biochemistry and Behavior* **24**: 1471-1477.
- Thiébot, M. H., Hamon, M. & Soubrié, P. (1982). Attenuation of induced-anxiety in rats by chlordiazepoxide: role of raphe dorsalis benzodiazepine binding sites and serotonergic neurons. *Neuroscience* **7**: 2287-2294.
- Thiébot, M. H., Le Bihan, C., Soubrié, P. & Simon, P. (1985). Benzodiazepines reduce the tolerance to reward delay in rats. *Psychopharmacology* **86**: 147-152.
- Thompson, R. F., Swain, R., Clark, R. & Shinkman, P. (2000). Intracerebellar conditioning — Brogden and Gantt revisited. *Behavioural Brain Research* **110**: 3-11.
- Thorndike, E. L. (1911). *Animal intelligence: experimental studies*, Macmillan, New York.
- Tiffany, S. T. & Drobos, D. J. (1990). Imagery and smoking urges: the manipulation of affective content. *Addictive Behaviors* **15**: 531-539.
- Tinbergen, N. (1948). Social releasers and the experimental method required for their study. *Wilson Bulletin* **60**: 6-52.
- Tolman, E. C. (1932). *Purposive Behavior in Animals and Men*, Century, New York.
- Tolman, E. C. (1934). Theories of learning. In *Comparative Psychology* (Moss, F. A., ed.), pp. 367-408. Prentice-Hall, New York.
- Tomie, A. (1996). Locating reward cue at response manipulandum (CAM) induces symptoms of drug abuse. *Neuroscience and Biobehavioral Reviews* **20**: 505-535.
- Tomie, A., Aguado, A. S., Pohorecky, L. A. & Benjamin, D. (1998). Ethanol induces impulsive-like responding in a delay-of-reward operant choice procedure: impulsivity predicts autoshaping. *Psychopharmacology* **139**: 376-382.
- Tomie, A., Aguado, A. S., Pohorecky, L. A. & Benjamin, D. (2000). Individual differences in Pavlovian autoshaping of lever pressing in rats predict stress-induced corticosterone release and mesolimbic levels of monoamines. *Pharmacology, Biochemistry and Behavior* **65**: 509-517.
- Trapold, M. A. (1970). Are expectancies based upon different positive reinforcing events discriminably different? *Learning & Motivation* **1**: 129-140.
- Trapold, M. A. & Overmier, J. B. (1972). The second learning process in instrumental learning. In *Classical conditioning II: Current research and theory* (Black, A. H. & Prokasy, W. F., eds.), pp. 427-452. Appleton-Century-Crofts, New York.
- Turken, A. U. & Swick, D. (1999). Response selection in the human anterior cingulate cortex. *Nature Neuroscience* **2**: 920-924.
- Tzschenkte, T. M. (2000). The medial prefrontal cortex as a part of the brain reward system. *Amino Acids* **19**: 211-219.
- Uylings, H. B. M. & van Eden, C. G. (1990). Qualitative and quantitative comparisons of the prefrontal cortex in rat and in primates, including humans. *Progress in Brain Research* **85**: 147-166.

- Vogt, B. A. (1993). Structural organization of cingulate cortex: areas, neurons, and somatodendritic transmitter receptors. In *Neurobiology of cingulate cortex and limbic thalamus: a comprehensive handbook* (Vogt, B. A. & Gabriel, M., eds.), pp. 19-70. Birkhäuser, Boston, Massachusetts.
- Vogt, B. A., Finch, D. M. & Olson, C. R. (1992). Functional heterogeneity in cingulate cortex: the anterior executive and posterior evaluative regions. *Cerebral Cortex* **2**: 435-443.
- Vogt, B. A. & Peters, A. (1981). Form and distribution of neurons in rat cingulate cortex: areas 32, 24, and 29. *Journal of Comparative Neurology* **195**: 603-625.
- Volkow, N. D., Ding, Y. S., Fowler, J. S. & Wang, G. J. (1996). Cocaine addiction: hypothesis derived from imaging studies with PET. *Journal of Addictive Diseases* **15**: 55-71.
- Volkow, N. D., Wang, G. J. & Fowler, J. S. (1997). Imaging studies of cocaine in the human brain and studies of the cocaine addict. *Annals of the New York Academy of Sciences* **820**: 41-54; discussion 54-45.
- von Neumann, J. & Morgenstern, O. (1947). *Theory of games and economic behavior*, Princeton University Press, Princeton, New Jersey.
- Wade, T. R., de Wit, H. & Richards, J. B. (2000). Effects of dopaminergic drugs on delayed reward as a measure of impulsive behavior in rats. *Psychopharmacology* **150**: 90-101.
- Wagner, A. R. (1978). Expectancies and the priming of STM. In *Cognitive Processes in Animal Behavior* (Hulse, S. H., Fowler, H. & Honig, W. K., eds.), pp. 177-210. Lawrence Erlbaum, Hillsdale, New Jersey.
- Wagner, R. F. (1970). Secondary emotional reactions in children with learning disabilities. *Mental Hygiene* **54**: 577-579.
- Walker, D. L. & Davis, M. (1997). Double dissociation between the involvement of the bed nucleus of the stria terminalis and the central nucleus of the amygdala in startle increases produced by conditioned versus unconditioned fear. *Journal of Neuroscience* **17**: 9375-9383.
- Warburton, E. C., Aggleton, J. P. & Muir, J. L. (1998). Comparing the effects of selective cingulate cortex lesions and cingulum bundle lesions on water maze performance by rats. *European Journal of Neuroscience* **10**: 622-634.
- Weinberger, N. M. (1995). Retuning the brain by fear conditioning. In *The Cognitive Neurosciences* (Gazzaniga, M., ed.), pp. 1071-1089. MIT Press, Cambridge, Massachusetts.
- Weinberger, N. M. (1998a). Physiological memory in primary auditory cortex: characteristics and mechanisms. *Neurobiology of Learning and Memory* **70**: 226-251.
- Weinberger, N. M. (1998b). Tuning the brain by learning and by stimulation of the nucleus basalis. *Trends in Cognitive Sciences* **2**: 271-273.
- Weissenborn, R., Robbins, T. W. & Everitt, B. J. (1997). Effects of medial prefrontal or anterior cingulate cortex lesions on responding for cocaine under fixed-ratio and second-order schedules of reinforcement in rats. *Psychopharmacology* **134**: 242-257.
- Whalen, P. J., Bush, G., McNally, R. J., Wilhelm, S., McInerney, S. C., Jenike, M. A. & Rauch, S. L. (1998). The emotional counting Stroop paradigm: a functional magnetic resonance imaging probe of the anterior cingulate affective division. *Biological Psychiatry* **44**: 1219-1228.
- Whishaw, I. Q. & Kornelsen, R. A. (1993). Two types of motivation revealed by ibotenic acid nucleus accumbens lesions: dissociation of food carrying and hoarding and the role of primary and incentive motivation. *Behavioural Brain Research* **55**: 283-295.
- White, K. G. & Pipe, M. E. (1987). Sensitivity to reinforcer duration in a self-control procedure. *Journal of the Experimental Analysis of Behavior* **48**: 235-250.
- White, N. M. (1997). Mnemonic functions of the basal ganglia. *Current Opinion in Neurobiology* **7**: 164-169.
- Whitelaw, R. B., Markou, A., Robbins, T. W. & Everitt, B. J. (1996). Excitotoxic lesions of the basolateral amygdala impair the acquisition of cocaine-seeking behaviour under a second-order schedule of reinforcement. *Psychopharmacology* **127**: 213-224.
- Wickens, J. & Kötter, R. (1995). Cellular models of reinforcement. In *Models of Information Processing in the Basal Ganglia* (Houk, J. C., Davis, J. L. & Beiser, D. G., eds.), pp. 187-214. MIT Press, Cambridge, Massachusetts / London.
- Wickens, J. R., Begg, A. J. & Arbutnot, G. W. (1996). Dopamine reverses the depression of rat corticostriatal synapses which normally follows high-frequency stimulation of cortex in vitro. *Neuroscience* **70**: 1-5.
- Williams, B. A. (1991). Marking and bridging versus conditioned reinforcement. *Animal Learning & Behavior* **19**: 264-269.
- Williams, B. A. (1994a). Conditioned reinforcement: neglected or outmoded explanatory construct? *Psychonomic Bulletin & Review* **1**: 457-475.
- Williams, B. A. (1994b). Reinforcement and choice. In *Animal Learning and Cognition* (Mackintosh, N. J., ed.), pp. 81-108. Academic Press.

- Williams, B. A. & Dunn, R. (1991). Preference for conditioned reinforcement. *Journal of the Experimental Analysis of Behavior* **55**: 37-46.
- Williams, D. R. & Williams, H. (1969). Auto-maintenance in the pigeon: sustained pecking despite contingent nonreinforcement. *Journal of the Experimental Analysis of Behavior* **12**: 511-520.
- Wilson, C. J. (1995). The contribution of cortical neurons to the firing pattern of striatal spiny neurons. In *Models of Information Processing in the Basal Ganglia* (Houk, J. C., Davis, J. L. & Beiser, D. G., eds.), pp. 29-50. MIT Press, Cambridge, Massachusetts / London.
- Wilson, P. N., Boumphrey, P. & Pearce, J. M. (1992). Restoration of the orienting response to a light by a change in its predictive accuracy. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **44**: 17-36.
- Wise, R. A. (1981). Brain dopamine and reward. In *Theory in Psychopharmacology Volume 1* (Cooper, S. J., ed.), pp. 103-122. Academic Press, London.
- Wise, R. A. (1982). Neuroleptics and operant behavior: the anhedonia hypothesis. *Behavioral and Brain Sciences* **5**: 39-87.
- Wise, R. A. (1985). The anhedonia hypothesis: Mark III. *Behavioral and Brain Sciences* **8**: 178-186.
- Wise, R. A. (1994). A brief history of the anhedonia hypothesis. In *Appetite: Neural and Behavioural Bases* (Legg, C. R. & Booth, D., eds.), pp. 243-263. Oxford University Press, New York.
- Wise, S. P. (1996). The role of the basal ganglia in procedural memory. *Seminars in the Neurosciences* **8**: 39-46.
- Wise, S. P., Murray, E. A. & Gerfen, C. R. (1996). The frontal cortex-basal ganglia system in primates. *Critical Reviews in Neurobiology* **10**: 317-356.
- Wogar, M. A., Bradshaw, C. M. & Szabadi, E. (1992). Choice between delayed reinforcers in an adjusting-delay schedule - the effects of absolute reinforcer size and deprivation level. *Quarterly Journal of Experimental Psychology, Section B - Comparative and Physiological Psychology* **45B**: 1-13.
- Wogar, M. A., Bradshaw, C. M. & Szabadi, E. (1993a). Does the effect of central 5-hydroxytryptamine depletion on timing depend on motivational change? *Psychopharmacology* **112**: 86-92.
- Wogar, M. A., Bradshaw, C. M. & Szabadi, E. (1993b). Effect of lesions of the ascending 5-hydroxytryptaminergic pathways on choice between delayed reinforcers. *Psychopharmacology* **111**: 239-243.
- Wolterink, G., Phillips, G., Cador, M., Donselaar-Wolterink, I., Robbins, T. W. & Everitt, B. J. (1993). Relative roles of ventral striatal D1 and D2 dopamine receptors in responding with conditioned reinforcement. *Psychopharmacology* **110**: 355-364.
- Wultz, B., Sagvolden, T., Moser, E. I. & Moser, M. B. (1990). The spontaneously hypertensive rat as an animal model of attention-deficit hyperactivity disorder: effects of methylphenidate on exploratory behavior. *Behavioral and Neural Biology* **53**: 88-102.
- Wyvell, C. L. & Berridge, K. C. (2000). Intra-accumbens amphetamine increases the conditioned incentive salience of sucrose reward: enhancement of reward "wanting" without enhanced "liking" or response reinforcement. *Journal of Neuroscience* **20**: 8122-8130.
- Zaborszky, L., Alheid, G. F., Beinfeld, M. C., Eiden, L.E., Heimer, L. & Palkovits, M. (1985). Cholecystokinin innervation of the ventral striatum - a morphological and radioimmunological study. *Neuroscience* **14**: 427.
- Zahm, D. S. & Brog, J. S. (1992). On the significance of subterritories in the "accumbens" part of the rat ventral striatum. *Neuroscience* **50**: 751-767.
- Zahm, D. S., Jensen, S. L., Williams, E. S. & Martin, J. R. (1999). Direct comparison of projections from the central amygdaloid region and nucleus accumbens shell. *European Journal of Neuroscience* **11**: 1119-1126.
- Zilles, K. (1985). *The cortex of the rat: a stereotaxic atlas*. First edition, Springer-Verlag, Berlin.
- Zilles, K. & Wree, A. (1985). Cortex: areal and laminar structure. In *The Rat Nervous System* (Paxinos, G., ed.), Vol. 1, pp. 375-415. Academic Press, Sydney.
- Zilles, K. & Wree, A. (1995). Cortex: areal and laminar structure. In *The Rat Nervous System*, Second edition (Paxinos, G., ed.), pp. 649-685. Academic Press, London.
- Zoli, M., Torri, C., Ferrari, R., Jansson, A., Zini, I., Fuxe, K. & Agnati, L. F. (1998). The emergence of the volume transmission concept. *Brain Research Reviews* **26**: 136-147.