

**CONTROL IN LIVING SYSTEMS: AN EXPLORATION OF THE
CYBERNETIC PROPERTIES OF INTERACTIVE BEHAVIOUR**

HEATHER C. BELL

Master of Science, University of Lethbridge, 2008

A Thesis

Submitted to the School of Graduate Studies
of the University of Lethbridge
in Partial Fulfillment of the
Requirements for the Degree

DOCTOR OF PHILOSOPHY
(Evolution & Behaviour)

Department of Neuroscience
University of Lethbridge
LETHBRIDGE, ALBERTA, CANADA

© Heather C. Bell, 2013

Dedication

In memory of William T. Powers (1926-2013), whose writings on living control systems were the initial inspiration for this work.

Abstract

Many types of animal behaviour, especially seemingly complex social interactions, have been attributed to the existence of complex cognitive mechanisms, underpinned by stimulus-response (S-R) rules. Indeed, as specific behaviours are analyzed in greater and greater detail, the increasing number of minor variations observed, even under tightly-controlled experimental conditions, seem to necessitate the operation of increasingly powerful computational devices. An alternate view, inspired by cybernetic theory, is that what is important is not the *specific* behaviours used by animals, but the *goal* of the organism in a particular context. In this thesis, a closed-loop cybernetic methodology for understanding behaviour is developed and implemented. Evidence is presented that, not only do at least some behaviours of animals function like engineered control systems, but also that this type of architecture is widespread in phylogenetic terms, relatively robust to interference, and able to be artificially reproduced. Implications for the study of the behaviour of all organisms are discussed.

Acknowledgments

Many thanks to everyone that assisted with the research: Erik Johnson, Brendan Long, and Lewis Horwood for hours of video analysis of a variety of creatures; Kevin Judge for help with the crickets and cockroaches; Evelyn Field and David Logue for providing video footage of dodging rats and fighting cockroaches, respectively; and Jeff Schank and Greg Bell for assistance with the development of the simulation.

I would also like to acknowledge my committee members, Dr. Louise Barrett, Dr. Michael Stingl, and Dr. John Vokey. Your thoughtful questions and input have kept me focused, and have driven me to think about my research on multiple levels, from a variety of perspectives, which has been invaluable. I would also like to thank Dr. Warren Mansell for agreeing to be my external examiner. Your interest in and encouraging comments on my research as it has progressed have always been welcome.

A gigantic thank-you to my exceptional friends and family, and especially to Frédéric. All of you have worked hard to keep me (relatively) sane. Your continuing support – despite what must seem, at the very least, to be a bizarre and erratic life trajectory – is baffling, but extremely appreciated, and I love you all.

Finally, I would like, most of all, to acknowledge my supervisor, Dr. Sergio Pellis. Your years of wisdom, encouragement, and guidance; of thought-provoking conversation on a multiplicity of topics; of level-headed temperance when I needed it; of books, food, wine, and friendship, have meant a great deal to me. I cannot imagine having had a better mentor during my graduate career. Thank-you, and I look forward to hopefully many future collaborations!

Disclosure of Publication

Some of the work included in this thesis has previously been published. The bulk of Chapter 2 and Chapter 3 appeared in *Animal Behaviour* (H. C. Bell & Pellis, 2011; H. C. Bell, Judge, Johnson, Cade, & Pellis, 2012). Additionally, large parts of the introduction and conclusion were included in a solo-authored manuscript that was recently submitted to *International Journal of Comparative Psychology* (H. C. Bell, submitted).

Contents

Approval/Signature Page	ii
Dedication	iii
Abstract	iv
Acknowledgments	v
Disclosure	vi
Table of Contents	vii
List of Tables	x
List of Figures	xi
1 INTRODUCTION	1
The S-R Model: A Very Brief History	3
The Variability of Behaviour: An S-R Conundrum	5
Cybernetics: An Alternative to S-R	6
Are Organisms Cybernetic? Some Clues From Behaviour	12
Developing the Methodology to Test the Cybernetic Properties of Behaviour	13
2 A CYBERNETIC PERSPECTIVE ON FOOD PROTECTION IN RATS: SIMPLE RULES CAN GENERATE COMPLEX AND ADAPT-ABLE BEHAVIOUR	15
Method	18
Subjects	18
Testing	18
Analysis	19
Results	20
Correlations	23
Variability	26
Individual Differences and Contextual Information	26
Discussion	31
3 HOW IS A CRICKET LIKE A RAT? INSIGHTS FROM THE APPLICATION OF CYBERNETICS TO EVASIVE FOOD PROTECTIVE BEHAVIOUR	35
Method	38
Subjects	38

Testing	39
Video Analysis	40
Results	44
Are Crickets Cybernetic Agents?	44
Are Crickets the Same as Rats?	48
Is Defender Orientation a Controlled Variable?	51
Does Approach Orientation Determine the Type of Evasive Strategy Used?	51
Discussion	58
Cybernetic versus Ballistic Movements	58
Commonalities in the Defensive Behaviour of Crickets and Rats	58
The Importance of Orientation	60
4 SIMULATING THE SYSTEM	65
Method	66
Properties of the Simulation	66
Analysis	68
Results	70
Correlations	70
Variability	70
Discussion	72
5 PERTURBING THE SYSTEM: EXPLORING THE PARAMETERS OF THE DODGING SYSTEM IN RATS	75
Experiment 1 and 2: General Method	79
Experiment 1: Disrupting Sensory Input I	80
Method	80
Analysis	80
Experiment 2: Disrupting Sensory Input II	81
Method	81
Analysis	81
Experiment 3: Disrupting Information Processing	82
Method	82
Analysis	83
Results	84
Correlations	84
Variability	84
Orientation	87
Differences Between Conditions	87
Discussion	90

6 CONCLUSION	94
Is Dodging in Rats Cybernetic?	94
How widespread is the cybernetic dodging rule?	96
Wiener’s approach: Can we simulate the behaviour?	98
How robust is the cybernetic system?	99
What if the Controlled Variable is Unknown?	100
Where to Go From Here? The Role of Neuromorphic Engineering	102
The Role of Embodied Cognition in the Production of Behaviour	103
How is Variable Behaviour Generated	104
General Conclusion	107
References	109
APPENDIX	122

List of Tables

2.1	Predictions and results.	22
2.2	Correlations for individual animals.	25
2.3	Coefficients of Variation	27
2.4	Partial and Semi-Partial Correlations	32
3.1	Competing Hypotheses and Predictions Tested	46
3.2	Coefficients of Variation, $\frac{\sigma}{\mu}$	50
3.3	Pattern of Correlations Among Variables for Crickets and Rats	52
3.4	Counts of Final Defender Orientations with respect to Robbers	54
3.5	Counts of Initial Robber Orientations with respect to Defenders	56
3.6	Correlations by Approach Orientation	57
4.1	Coefficients of Variation, $\frac{\sigma}{\mu}$	73
5.1	Competing Predictions with respect to Three Hypotheses	78
5.2	Coefficients of Variation, $\frac{\sigma}{\mu}$	88
5.3	Frequencies of Final Defender Orientation by Experimental Group and Corresponding χ^2 Values	89
5.4	Overall Changes to Dodging Parameters, Relative to Controls	91

List of Figures

2.1	Illustration of dodging and measures used. Columns A) and B) depict the same dodge with either the distance, or the angle measurements indicated. a) is the frame during which the dodge was initiated, b) is the frame during which the robber ceased robbing, and c) is the frame during which the dodger ceased dodging. A)a. corresponds to d_i , A)b. shows d_r , and A)c. depicts d_f , as measured using the greyed-out “virtual” location of the robber when robbing ceased (even though robbing actually ceased during a previous frame.) B)b. shows <i>robber angle</i> , and B)c. shows <i>dodger angle</i>	21
2.2	Correlations between a) robber angle and dodger angle, and b) robber angle and inter-animal distance.	24
2.3	Mean inter-animal distance by subject at three timepoints: A) initial inter-animal distance, d_i ; B) inter-animal distance when robber stopped robbing, d_r ; C) inter-animal distance when dodger stopped dodging, d_f . Error bars are 95% confidence intervals.	29
2.4	Initial inter-animal distance, d_i , as a function of food type. Error bars are 95% confidence intervals.	30
3.1	Illustration of how the robber orientation was determined. D is the defender. R1 is a robber in the front orientation; R2 is a robber in the side orientation; R3 is a robber in the rear orientation. The horizontal lines denote the cut-off criteria used.	45
3.2	A. Variability in path length in running crickets over 4 interactions. B. Variability in path length in dodging crickets over 4 interactions. Whiskers indicate range, boxes are first and third quartiles, and centre bar is the median.	47
3.3	A. Correlation between robber path length and defender path length. B. Correlation between robber path length and inter-animal distance.	49
3.4	A. Mean inter-animal distance for running crickets at three time points: a) d_i ; b) d_r ; and c) d_f . B. Mean inter-animal distance for the same three time points as in A. for dodging crickets.	53
4.1	Simulated rats behaving in a circular arena. 1 is the position of the defender before the robber approaches, and 2 is the position of the defender after the approach of the robber.	69
4.2	A. Correlation between robber path length and defender path length. B. Correlation between robber path length and inter-animal distance.	71
5.1	Robber path length vs. inter-animal distance for A. Controls, B. Vibrissae Removed, C. Vibrissae Removed and Enucleated, and D. Isolates	85

5.2	Robber path length vs. Defender path length for A. Controls, B. Vibrissae Removed, C. Vibrissae Removed and Enucleated, and D. Isolates	86
6.1	Four division of the cockroach body	124
6.2	Proportion of contacts to each quarter	126
6.3	Proportion of flips relative to contacts on each quarter	127

Chapter 1

INTRODUCTION

The vast majority of ethological, psychological, and neuroscience (EPN) research, grounded in Behaviourist and Cognitive Theory, assumes that behaviour is fundamentally stimulus-response (S-R) – that is, a behavioural response is the result of exposure to a stimulus, either external or internal. In its modern instantiations, specific behaviours are thought to be learned by an organism as a means of controlling its environment (e.g., Rescorla, 1988). For example, a rat learns, through an understanding of the consequences of its actions, to press a bar because it understands that the outcome of this action is rewarding, in that it obtains a food pellet. Although the details of how responses are elicited vary in this example, depending on the specific theory used to explain it (e.g., Dickinson & Balleine, 2000; Sutton & Barto, 1998), these theories all contain the same underlying assumption that there is a linear and unidirectional relationship between stimulus “input” and behavioural “output” – even if it is the consequence of the behavioural output, not the behaviour itself, that is important to the organism. If this is true, then if the conditions under which a rat learns to press a bar are exactly the same, it should always press the bar using the same behaviour – that is, its behaviour should be completely predictable. However, even in tightly-controlled experiments, the behaviour of organisms is variable (Breland & Breland, 1961; Neuringer, 2004). And in fact, it may be the case the behaviour is actually fundamentally variable (Blough, 1966; Korobkova, Emonet, Vilar, Shimizu, & Cluzel, 2004; Neuringer, 2002; Park, Pontius, Guet, Emonet, & Cluzel, 2010). If behavioural variability is fundamental, one would expect to observe what we do – broadly similar behavioural patterns, with behavioural variability playing little role, except as noise (however, see Brembs, 2011). However, models like those described

above typically account for less than 40% of the observed variance in behaviour (and often much less) (e.g., Armitage & Conner, 2001).

There is, however, an alternative to the above view, based on the ideas of cybernetics (Rosenblueth, Wiener, & Bigelow, 1943; Rosenblueth & Wiener, 1945; Wiener, 1948). If we re-conceptualize behaviour in different terms, behavioural variability takes on a more central role – becoming a necessary component of behaviour, instead of unexplained noise. Rather than an organism controlling the outcome of an event by learning a specific behaviour through its consequences (reward, negative reinforcement, or punishment), instead what the organism might be doing is varying its behaviour in real time in order to control the levels of relevant variables. In the case where food is used as a reinforcer for a bar press, bar pressing itself becomes more likely, because bar pressing – however that is actually accomplished – ultimately reduces hunger. Hunger level, not reward, is what is being controlled by the animal. And in order to control hunger, the organism must first learn to control other environmental variables that are related to the control of hunger, such as the orientation of its body with respect to the bar. Unlike the control of a behavioural outcome, control in this sense is an ongoing, dynamic process. Variable behaviour is required for organisms because, unlike the environment inside an operant chamber, the real world, in which organisms evolved to navigate, is not static. Any number of minor to major disturbances, which are often unpredictable, can prevent an organism from, for instance, controlling its hunger level – or any of the perceptual variables related to hunger. Variability provides a means of compensating for these disturbances. This may seem like a minor and obtuse re-interpretation of the bar-pressing phenomenon, but it has major implications for the understanding of behaviour. And crucially, models of this sort can account for as much as 99% of the variance observed in behaviour (e.g., Marken, 1986, 1990).

Despite the fact that the formal field of cybernetics has existed for some time (Rosenblueth et al., 1943; Rosenblueth & Wiener, 1945; Wiener, 1948), the ideas generated by it have largely failed to be adopted, in an empirical sense, by researchers in EPN – even though much work in artificial intelligence and robotics assumes that the behaviour of organisms is cybernetic. In this thesis, it is argued that, not only does a cybernetic model provide a more powerful explanatory platform for understanding behaviour than do linear models, but that variability is a crucial component of behaviour. And, as opposed to linear models, which lead researchers to focus on what an animal is doing at a specific point in time in relation to a set of inputs, a cybernetic approach views behaviour in terms of controlled variables – that is, the aspect(s) of the organism’s internal or external environment that the behaviour of the organism is produced in the service of changing or maintaining. In addition, although it is not antithetical to linear models of behaviour, the idea that behaviour is constrained by physics is discussed. Further, it is argued that the rules that animals use to behave, cybernetic or not, are not abstract concepts encoded solely in the brain, but embedded in a larger system that includes the organism and its environment, which is an important consideration for a complete understanding of behaviour.

The S-R Model: A Very Brief History

In its infancy, Psychology strove to establish itself as a respectable science. Because the accepted model of experimental science at the time was derived from Newton’s classical view of physics – that is, the linear sum of the external forces acting on an object fully explains the behaviour of that object Psychology adopted this approach to understanding the behaviour of organisms as well, despite the fact that early researchers, such as Wundt and James, both recognized the role of purposefulness,

and thus, an ability to be not only influenced by, but also to influence the environment, in the behaviour of living things (Cziko, 2000).

Gradually, it became clear that some of the early psychological methodologies, such as the use of introspection, lacked the strict objectivity of proper scientific methods. Beginning with Pavlov, whose ideas were further developed by Watson and Skinner, a new paradigm in Psychology emerged – one that emphasized understanding the inputs into the system (the stimuli to which the organism was exposed) in relation to the resulting output (the behavioural response of the organism). Although Behaviourism offered some clear advantages over introspection in terms of standardization of experiments, replication, control, and objectivity, at its heart, it still relied on Newton’s linear causality model.

Eventually, in the course of the Behaviourist program, it became evident that not every animal could be reliably trained to perform any task, and that even under highly controlled conditions, behaviour was variable (Breland & Breland, 1961). This violated the behaviourist assumption that the behaviour of organisms is simply sums of their inputs. It was recognized that, at the very least, phylogeny is also a determining force in the production of behaviour.

In attempts to address these issues and others, and in conjunction with ideas emerging from the development of the computer, the cognitive revolution was born. Unlike behaviourism, cognitive approaches asked questions about the mechanisms of how input was processed by the machinery of the brain to create output. Effectively, cognitive theory adds one additional element to the behaviourist S-R model: an intermediate processing step, S-O-R, where the O stands for the organism (Hebb, 1949). Although many cognitive theories include various sorts of feedback mechanisms within the processing aspect, they are still fundamentally linear theories. Similarly, early ethological theories posited complex internal mechanisms (e.g., the innate releasing

mechanism), and so had more in common with modern cognitive approaches than they did with contemporary behaviourist theories. Nonetheless, like behaviourism and cognitivism, in ethology the animal's behaviour is in large part dependent on the incoming input, as particular 'fixed action patterns' (phylogenetically shaped motor outputs) are released by appropriate sign stimuli (Lorenz, 1981). That is, ethological, behaviourist and cognitive theories, which underpin the majority of current EPN research, all ignore the fact that, not only does the environment act on organisms, but that organisms also act on the environment. An inevitable consequence of the S-R view is that, at the very least, important sources of behavioural variability are missed.

The Variability of Behaviour: An S-R Conundrum

Under the current major theories of behaviour, behavioural variability leads to the conclusion that all inputs have not been fully determined nor understood. It would seem, therefore, that precise behavioural prediction is, at least for practical purposes, impossible.

Behavioural variability can refer to either within-organism variability in the same context, or between-organism variability in the same context. Between-organism variability is generally seen as either genetic in nature or due to some early developmental experience – some fundamental preference – whereas within-organism variability is seen to be contextually dependent – the sum of the inputs to the organism at a particular time. Between-organism variability is well-known, and is the reason that within-subject designs tend to be more powerful – and one might suspect that small phenotypic variation could account for these individual differences. However, even in genetic clones reared in the same environments, behaviour is highly variable (e.g.,

Sondergaard, Herskin, Ladewig, Holm, & Dagnaes-Hansen, 2012). Within-organism variability, seen in even tightly controlled experiments, presents an even greater problem in terms of explanation.

Although behavioural variability is well known in “higher” animals, such as mammals, it is, to some degree, acceptable, given the complexity of their nervous systems. However, even the simplest of animals exhibit variability. For example, fruit flies (*Drosophila melanogaster*) are highly variable in their tendency to turn left or right in the optomotor task, and mathematical analysis suggests that this variability is not simply the result of noise (Maye, Hsieh, Sugihara, & Brembs, 2007). Paramecia are known to switch direction spontaneously, even in homogeneous environments (Oosawa & Nakaoka, 1977). And even the lowly bacterium, *Escherichia coli*, exhibits what appears to be inherent behavioural variability that cannot be explained by appealing to the sum of its inputs (Korobkova et al., 2004; Park et al., 2010).

One could argue that the problem is simply that all of the inputs to the system are not known, or that they are not known precisely enough. However, even in very tightly controlled behavioural experiments, where the conditions between trials differ as little as possible, behaviour seems to vary more than should be accountable for by appealing to the few underlying input variables whose values remain in flux or unknown (e.g., temporal input).

Cybernetics: An Alternative to S-R

The problem of behavioural prediction in a proximate sense – that is, what the organism will do, specifically, from moment to moment – is difficult to resolve, in light of behavioural variability. Will a cat bat you with her paw, or meow loudly in your ear to get your attention? What if this is not the important question, or at least

not the whole question? What if we ask, instead, what is the goal of the organism – what are the important variables that the organism might be trying to control with its behaviour?

The word *cybernetics* is derived from the Greek term for the steersman of a ship, describing a self-correcting system that maintains a goal by taking compensatory actions to offset the effects of disturbances. The field of cybernetics borrows from the physiological notion of homeostasis (i.e., variables being maintained via negative feedback systems, or cybernetic *control structures*). Norbert Wiener and his colleagues were the first to explicitly recognize that the logic behind self-regulating systems could apply equally well to the overt behaviour of living things (Wiener, 1948), in which the animal varies its behaviour in order to maintain some aspect of its relationship to the environment constant. William Powers built on this by proposing an explicitly cybernetic theory specifically designed to explain behavioural and psychological phenomena, Perceptual Control Theory (PCT) (Powers, 1973). The basic tenet of cybernetic theory is that *goal states* are achieved by organisms through the implementation of negative feedback control mechanisms. Importantly, cybernetic theory acknowledges that organisms act on their environments. In this view, behaviour is a continuous process of adjustment by the organism as it attempts to reduce the error (difference) between what it perceives and its goal state. A goal does not necessarily imply conscious purposefulness on the part of an organism. Any homeostatic system has a preferred (goal) state, which is the set point of the parameter of interest for that system. The system does not need to explicitly know what the goal state is.

Cybernetic models provide two practical features of direct relevance to the study of animate behaviour. First, that what animals maintain constant is some perception that they have of the world, and that behaviour is thus part of the means to do so. For example, when driving, we keep the side of the front of the car oriented a certain

distance to the centerline. Potholes, gusts of wind and other cars may disturb that relationship, leading to movements of the steering wheel (by our hands, and hence our behaviour) to regain and maintain the perceptual relationship between the car and the centerline. The behaviour is variable and that variability is in the service of maintaining a constant perception.

Second, although negative feedback control structures comprise the smallest functional units within cybernetic models, these units can be arranged hierarchically, thus making it possible for organisms to adapt their behaviour to maintain multiple controlled variables (CVs) (i.e., organisms can have multiple “goals”), and for there to be an order of precedence for achieving these goals.

A classic example of a non-biological control system is a household thermostat. Its “goal” is to maintain a particular temperature that has been set by the operator. If the temperature that the thermostat measures is lower than the goal temperature, it activates the furnace. The activation of the furnace changes the temperature in the room, and hence the input to the thermostat, until the goal temperature is reached, at which point, the furnace is shut off. In essence, the behaviour of the thermostat itself affects its subsequent behaviour (negative feedback). Importantly, the thermostat does not “know” what the goal temperature is. And a thermostat has only two behavioural states – furnace on, or furnace off.

One could view the turning on and shutting off of the furnace by the thermostat as independent S-R events, noting that the furnace was turned on when the temperature was too low, and that it was shut off when the ambient temperature of the room matched the temperature set by the operator. But what was that actual cause of, in particular, the shutting off of the furnace at the matched temperature? In fact, it was the previous behaviour of the thermostat (turning on the furnace) that caused its own later behaviour (shutting off the furnace). Part of the input to the system was

the system's own behaviour. Artificially parsing the events into disparate temporal events causes the observer to miss the fundamental dynamics of the behaviour of the system as a whole –that is, that the two events are not unrelated, and in fact, are intrinsically interconnected.

Further if the owner of the house prefers a cooler ambient temperature when going to bed, the setting on the thermostat might be lowered. In this case, the owner is acting as a higher-level control system, affecting the functioning of the lower order system by changing the reference signal. Cybernetic models also allow for the modification of the specific reference values of CVs. Switching between homeostatic reference values (rheostasis) has been shown to occur in physiological systems, such as when body temperature is elevated during a fever (Mrosovsky, 1990). The hierarchical nature of cybernetic models solves two problems: first, it accounts for how homeostasis is maintained (constancy in the face of disturbance), and second, it accounts for how a different reference value can be achieved and defended.

Cybernetic approaches differ from other cognitive approaches to behaviour in that much of what is known about the structure and operation of organisms supports control structures interpretations. Homeostasis is a well-established phenomenon. At the most basic level, endogenously oscillating neurons, as well as half-centre (coupled) oscillators, and larger neural networks that function as oscillators, including those involved in memory, are thought to oscillate via negative feedback systems (Friesen & Block, 1984; Friesen, Block, & Hocker, 1993; Buzsáki, 2006). Larger networks occurring not just within the nervous system, such as the hypothalamic-pituitary-adrenal axis, also function to maintain homeostasis and rely on control structures to do so. It is not so far-fetched to postulate that the structure of the somatic nervous system, which is dynamically coupled to the outside world, would function the same way.

Several of the founders of cybernetics speculated, in one of the earliest papers on the subject, that many behaviours observed in organisms would turn out to be of the type found in artificial control structures such as thermostats (Rosenblueth et al., 1943), and Ashby (1952) recognized that the control of at least some variables, what he called “essential variables”, should be selected for by evolution. In later work, they outlined ways in which the control system properties of organisms could be tested using robotics (Rosenblueth & Wiener, 1945). Since that time, at least in the west, cybernetic theory has been applied to robotics and artificial intelligence (e.g., Johnston, 2008; Nahodil & Vitku, 2012), and has been fairly extensively used in sociology (Robinson, 2007), and has even been applied to clinical psychology (Carey, 2006). However, researchers seem to have generally skipped the investigation of the actual behaviour of the organisms in favour of moving straight to application, where it has often proved useful. Because of this, it is all the more puzzling that EPN has remained more or less firmly entrenched in S-R ideology (Marken & Mansell, 2013).

Others have proposed theories that have employed various types of feedforward and feedback in their architecture, such as Reinforcement Learning (Sutton & Barto, 1998), and the Associative Cybernetic Model (Dickinson & Balleine, 2000), as well as the work of Tolman (Tolman, 1932, 1948). On the surface, these seem compatible with PCT; however, there is one major difference. In these models, feedforward and feedback mechanisms are used to refine the acquisition of a specific behaviour, generally by modifying the organism’s prediction of the consequences of this behaviour. That is, feedback is not used continuously to modify behaviour itself, but rather to change the probability that a specific behaviour will re-occur, given a set of inputs. (Rescorla, 1988) has noted that modern behavioural theory explains behaviour as the method used by the organism to control its environment, by correcting the misalignment between a goal (consequence of behaviour) and its present reality. This

requires not only that an organism understand the consequences of its actions, but also that it has an internal representation of its ideal state. It is no wonder that many make a distinction between “automatic” S-R tasks and true goal-directed behaviour, attributing the latter only to more sophisticated organisms – and then, only in some situations (e.g., Dickinson & Balleine, 2000).

In contrast, a control system does not require an internal representation of the consequences of its actions, nor is there a true distinction between automatic and goal-directed behaviour. All behaviour is seen to be goal-directed, according to PCT, as it is performed in the service of controlling relevant variables. Control systems do not require internal representation because they only monitor the level of a perceptual variable and contrast that against the ideal level, producing variable behaviour if the error is too high. Since the error signal from lower-order control structures can be fed into higher-order structures (Vancouver, 2005), one can imagine that, given a complex enough system of control structures controlling a large enough number of variables, behaviour might appear as though the organism has an internal representation of the world and of the consequences of its actions. This basic idea is demonstrated in the simulation section of this thesis in a non-learning context.

In a learning context, modern behavioural theories (e.g., Dickinson & Balleine, 2000; Rescorla & Wagner, 1972; Sutton & Barto, 1998) describe the content of learning in terms of the acquisition of specific behaviours that minimize the disparity between the organism’s internal representation of its ideal state and its actual state, with respect to various stimulus inputs. Alternately, PCT describes learning as the process of altering the parameters of control systems that underlie behaviour, which is termed “re-organization.” In situations where there is little environmental disturbance and the conditions are often the same, high-level re-organization does not need to occur, and compensation need only take place at very low, sensory levels, so

the behaviour appears to be automatic or stereotyped. Although a full treatment of learning is outside the scope of this thesis, it is described in detail, including some potential neurophysiological underpinnings, by Yin (2013). Importantly, the principle of re-organization has been demonstrated by the creation of both simulated and robotic systems using hierarchically-arranged control structures that learn by randomly shifting control system parameters (e.g., Powers, 2008; Young, 2000). The resulting acquired behavioural patterns observed in the simulated entities are broadly similar in their execution over time, but exhibit the kinds of behavioural variability seen in real organisms, and crucially, are adaptable in the face of the introduction of various types of disturbances. This learning occurs, despite the fact that the simulated organisms lack an internal representation of any sort, other than the preferred value for the specific parameter with which they are concerned.

Although the theory and philosophy of cybernetics has been fairly well developed, even for the application to EPN (e.g., Ashby, 1956; Cziko, 2000; Powers, 1973; Wiener, 1948), and its usefulness for the construction of artificial systems has been established, the basic premise that organisms behave like control systems as opposed to S-R systems has remained, by and large, untested empirically.

Are Organisms Cybernetic? Some Clues From Behaviour

Although the majority of EPN research describes organisms as S-R machines, there is evidence in the literature that many organisms function like control systems. For example, to catch thrown objects, both dogs and humans maintain constant the linear optical trajectory – that is, the angle of the object on the retina with respect to the horizon (McBeath, Schaffer, & Kaiser, 1995; Schaffer & McBeath, 2002; Schaffer, Krauchunas, Eddy, & McBeath, 2004). Predatory flatworms track prey by maintain-

ing constant the perception of mucus trails (Iwai, Sugiura, & Chiba, 2010). Male crickets performing the 'judder', the back and forth movement performed in agonistic interactions, keep their bodies horizontal to the substrate, regardless of the incline of the substrate, by varying the specific movements used (Pellis, Gray, & Cade, 2009). Jellyfish appear to have a number of internal reference values for perceptual variables that they control with their behaviour using feedback systems: salinity levels, turbulence, somatosensory stimulation, the formation of aggregations, and direction of swimming (Albert, 2011). Spontaneous magnetic alignment in many types of organisms suggests that geomagnetic direction is a controlled variable (e.g., Nogueira & Lins de Barros, 1995; Rothsey & Rohde, 2002; Begall, Cervený, Neef, Vojtech, & Burda, 2008; Vácha¹, Kvíčalová, & Půžová, 2009). The carrion beetle, *Necrophorus humator*, maintains constant a certain wind direction with respect to the direction in which it is traveling (Böhm, Heinzl, Scharstein, & Wendler, 1991). In birds, although the specific movements used to accomplish flight vary widely, the aerodynamic forces involved are controlled in a consistent direction with respect to gravity (Dial, Jackson, & Segre, 2008). Complex social interactions in mammals and birds have also been shown to involve the homeostatic maintenance of inter-animal distance and/or inter-animal bodily orientations (e.g., Golani, 1976; Moran, Fentress, & Golani, 1981; Pellis, 1982). In all these cases, the animals are varying their behaviour in order to keep constant some perception (Powers, 1973).

Developing the Methodology to Test the Cybernetic Properties of Behaviour

Given the evidence that at least some behaviour might better be described in cybernetic, rather than S-R terms, the objective of the present work was to develop some

methodology 1) to test whether or not a particular behaviour is cybernetic, and 2) to explore the parameters of a cybernetic system.

A well-defined behaviour with an existing S-R description was chosen for study. In this thesis, a relatively simple methodology is developed, using robbing and dodging food protective behaviour in both rats and crickets, to test the cybernetic properties of this behaviour.

Chapter 2

A CYBERNETIC PERSPECTIVE ON FOOD PROTECTION IN RATS: SIMPLE RULES CAN GENERATE COMPLEX AND ADAPTABLE BEHAVIOUR

Robbing and dodging in rats involves one animal (the *dodger*) possessing a small piece of food, and another animal (the *robber*) attempting to acquire the food. The robber approaches the head of the dodger, and the dodger evades by swerving laterally away (e.g., Whishaw, 1988; Whishaw & Tomie, 1988; Whishaw, DuBois, & Field, 1998). This relatively simple behaviour offers an experimental paradigm suitable for studying brain mechanisms of motor control and inter-animal coordination (Whishaw & Oddie, 1989; Whishaw, Oddie, McNamara, Harris, & Perry, 1990; Field, Whishaw, & Pellis, 1996, 1997a, 1997b; Field, Whishaw, Forgie, & Pellis, 2004; Field, Watson, Whishaw, & Pellis, 2005; Pellis et al., 2006; Pellis, Field, & Whishaw, 1999). Furthermore, given the apparent correlation between the angle swept through by the dodger to evade the robber with both the type of food being consumed and the identity of the robber, robbing and dodging would appear to provide a relatively simple motor test for studying various aspects of cognition (Whishaw & Gorny, 1994; Pellis et al., 2006). However, viewing dodging as a simple lateral swerve, the magnitude of which is modulated by cognitive assessments of food value and other contextual factors, creates two explanatory difficulties.

The first problem is that, if the rats calculate the angular displacement appropriate for a given dodge, they would have to calculate the time it takes to eat a piece of food, then enter that information into an algorithm along with other variables, such as robber identity and movement. Moreover, the algorithm would require updating to incorporate adjustments needed due to the distance and approach speed of the robber. Such calculations posit the existence of cognitively demanding mechanisms

that would seem inefficient. Indeed, in opposition to an increased cognitive load, there is growing evidence that relatively simple cognitive processes or rules can generate complex behavioural outcomes (Pfeifer & Bongard, 2007; Barrett, 2011).

The second problem is that robbers rarely succeed in robbing (unpublished observations), yet rats can learn rapidly, so if a dodger were simply taking into account the time it takes to eat a piece of food and then calculating a fixed dodge angle, the robber would also soon learn what that dodge angle would be, and the robbing attempts should quickly become successful. Given that rats make dodging decisions and movements in fractions of a second and most often do so successfully, blocking the robber's access to the food item, some simpler rule must be being used, one that cannot be anticipated by the opponent. Given also that rats can judge inter-animal distance using visual and tactile cues (Pellis et al., 1996a), it was hypothesized that dodgers gain and maintain a constant inter-animal distance, with compensatory movements used to adjust the magnitude of the dodge to the movements of the robber. Any computational change with regard to food type/eating time could be accomplished by simply modifying the inter-animal distance maintained. This simple rule would eliminate the need to perform complex and time-consuming computations, and allow the dodger to evade robbing attempts successfully, regardless of the robber's strategy. The proposed rule involves only the maintenance of a constant distance between the two animals – as opposed to the calculation of a dodge angle based on contextual information – yet it explains dodging behaviour more satisfactorily (i.e., it can also account for learning on the part of the robber). In addition, fewer cognitive demands are made of the animals. Hence, the “maintain a constant distance” rule is more parsimonious than the “standard” S-R explanations.

The rule proposed is a cybernetic one, with which the defending rats gain and maintain a particular distance from the robber by varying their dodging behaviour.

For a variable to be considered *controlled* (in this case, inter-animal distance), it must be uncorrelated with system disturbances. In other words, the organism must be able to compensate for system disturbances in order to maintain some perceptual constancy. But one would expect that the movement of one animal would be correlated with the movement of the other animal, and confirming such a correlation is a method by which to test further that the presumed controlled variable is not simply an artefact of measurement error. That is, a variable, in this case, the distance between the dodger's nose and the robber's nose, is controlled by the dodger, and to do so, the dodger must vary its behaviour to stabilize the inter-animal distance (Powers, 1973).

If a dodger moves so as to achieve some stable distance between itself and the robber, then it is predicted that inter-animal distance (hereafter, *distance*) should be uncorrelated with system disturbances (the movement of the robber). Second, it is predicted that, because *dodger angle* should be a compensatory action in order to maintain the controlled variable, *distance*, that *dodger angle* (the angle traversed by the dodger during the dodge) should be correlated with the movement of the robber, *robber angle*. Third, it is predicted that, if *distance* is a controlled variable, *distance* should be significantly less variable than *dodger angle* – that is, the specific angle employed to evade the robber should be irrelevant, and thus, variable, whereas *distance* should be relatively stable. Fourth, it is predicted that, if *dodger angle* is irrelevant, *dodger angle* should not vary as a function of contextual information (e.g., qualities of the food being eaten, or of the robber). Fifth, because it is thought that the dodger is responding to the action of the robber, and not to contextual cues, it is expected that as the robber moves more quickly, that the dodger would need to increase the inter-animal distance in order to compensate; hence, it is predicted that the final distance between the two animals should be more related to the distance

between the animals when the robber stops robbing than it is to the distance between the animals when the dodge begins. In other words, we would predict that, in the course of the execution of the dodge, that the dodger should be using information about the behaviour of the robber in order to make dodging decisions, and whatever information was used at the initiation of the dodge, by the time the dodge is finished, should become less and less relevant. Finally, if individual rats are following a simple cybernetic rule, then all individuals should exhibit the attainment and maintenance of a specific inter-animal distance, even if there is variation across individuals in the actual distance defended. Therefore, it is predicted that each rat will follow the general predictions made above.

Method

Subjects

Twenty-one male Long-Evans Hooded rats derived from six litters, 40-60 days of age, that were born and raised at the University of Lethbridge, Lethbridge, Alberta, Canada, were used. Of the 21, seven animals were chosen to be subjects such that, when housed in triads, the triads were composed of the subject, plus one larger and one smaller male. Triads, including subjects, were housed in polyethylene tubs, on a 12h light/dark cycle with lights on at 0700h, and given water *ad libitum*.

Testing

A standard robbing and dodging paradigm was used (e.g., Field et al., 1996; Whishaw et al., 1998; Pellis et al., 2006). Subjects were individually habituated to a testing enclosure (a clear Plexiglas[®] cylinder, 40cm in diameter x 45cm high, with an inclined

mirror mounted underneath), until each would immediately approach and eat three pieces of food in succession. Prior to testing, the rats were food deprived to 80 - 85% of their free feeding body weight and were subsequently maintained on this food deprivation schedule until testing was completed (approximately 2 weeks in total). After each trial, the rats were weighed, and food was adjusted accordingly in order to maintain the desired 80 - 85% free feeding body weight.

Trials were conducted by placing the subject with one of its cagemates (larger or smaller than itself) in the cylindrical enclosure, and introducing one of two food types used. Foods used for the trials were raw azuki beans (Nutters Bulk & Natural Foods, Lethbridge, AB) and 300 mg rodent food pellets (banana pellets) (BioScience Inc., Frenchtown, NJ), which are of similar size, but differ in the amount of time required to consume. Trials were counterbalanced, and were filmed using a Sony 8mm camcorder at 30 frames per second.

Analysis

The video of the animals robbing and dodging was played back, and the experimenter selected 36 dodges for each subject. Because dodging against a wall is likely to influence dodge magnitude irrespective of other contextual information, only dodges that occurred in the centre of the enclosure were selected, and, as such, the first 36 dodges that met this criterion were used. In most cases, 36 dodges was the total number of dodges found for each subject in the video footage that met the above criteria. For one subject, only 34 dodges could be analyzed. The total number of dodges was made up of 9 dodges with the smaller animal in the *azuki* condition, 9 dodges with the smaller animal in the *banana* condition, 9 dodges with the larger animal in the *azuki* condition, and 9 dodges with the larger animal in the *banana*

condition.

Once the dodges had been selected, they were analyzed, frame-by-frame, using Vicon Motus motion capture software (Vicon Motion Systems, Colorado, USA). The positions of both the robber and the dodger – the tip of the nose, the nape of the neck, and the base of the tail – were digitally tagged and tracked throughout each video sequence. Following the digitizing of each video, Vicon Motus was used to calculate inter-animal distance measurements – the distance between the noses of the two animals – at all time points. *Distance*, the inter-animal distance for every timepoint, was used in some analyses. In addition, three *specific* timepoints were used for some analyses. These were the inter-animal distance when the dodge was initiated, d_i (Figure 2.1A.a.) and the inter-animal distance when the robber stopped robbing, d_r (Figure 2.1A.b.). At the third timepoint, inter-animal distance when the dodger stopped its lateral movement, d_f , was recorded as the distance between the robber when the robber stopped moving, and dodger when the dodger stopped dodging, even though these events occurred in different frames (see Figure 2.1A.c.). This was done to eliminate changes in inter-animal distance that arose from irrelevant movement by the robber (e.g., wandering off in a different direction) after the pursuance of the dodger had ended. In addition to the distance measurements, *robber angle* – the angle traversed by the robber (Figure 2.1B.b.) and the *dodger angle* – the angle traversed by the dodger (Figure 2.1B.c.) were also obtained using Vicon Motus and used in some analyses.

Results

The results are discussed in detail below. See Table 2.1 for a summary of the predictions made and the results obtained.

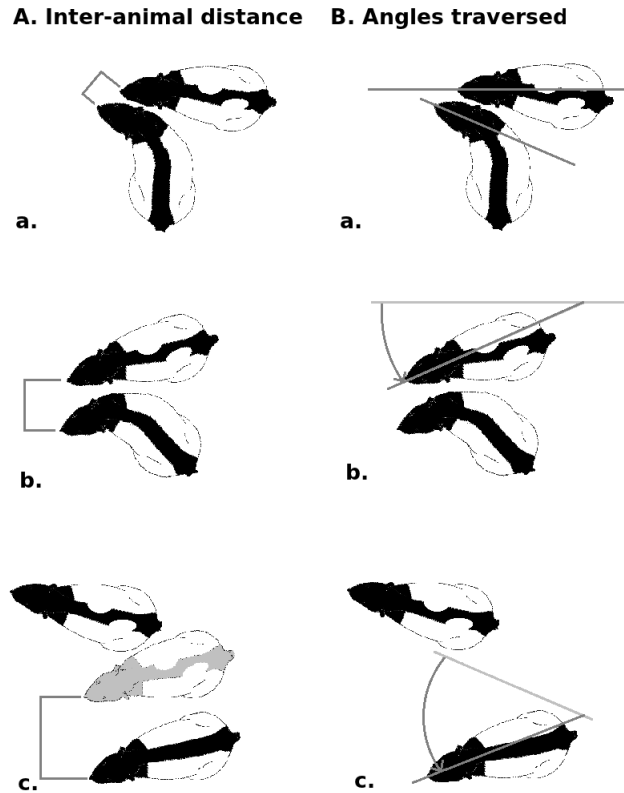


Figure 2.1: Illustration of dodging and measures used. Columns A) and B) depict the same dodge with either the distance, or the angle measurements indicated. a) is the frame during which the dodge was initiated, b) is the frame during which the robber ceased robbing, and c) is the frame during which the dodger ceased dodging. A)a. corresponds to d_i , A)b. shows d_r , and A)c. depicts d_f , as measured using the greyed-out “virtual” location of the robber when robbing ceased (even though robbing actually ceased during a previous frame.) B)b. shows *robber angle*, and B)c. shows *dodger angle*.

Table 2.1: Predictions and results.

Prediction	Obtained Result
Correlation between <i>distance</i> and <i>robber angle</i> should not be significant	Confirmed
Correlation between <i>robber angle</i> and <i>dodger angle</i> should be significant	Confirmed
Above pattern of correlations should hold for each animal	Confirmed
<i>Dodger angle</i> should be more variable than <i>distance</i>	Confirmed
Subjects should show individual differences in d_i	Confirmed
d_i and d_r can be modified by contextual information	Confirmed
d_i and d_r still depend <i>more</i> on robber behaviour than contextual information	Confirmed

Correlations

Two main predictions were made with respect to correlations between measures. First, it was expected that, if the distance between the two animals were a controlled variable, it should not be correlated with disturbances to the system. That is, if the robber were trying to move closer to the dodger to get the food (the system disturbance), but the dodger was able to compensate by varying its movements to keep inter-animal distance constant (the controlled variable), the value of the system disturbance variable should change, while the value of the controlled variable should remain stable. Hence, there should be no relationship between the system disturbance variable and the controlled variable. In this analysis, *robber angle* was used as a proxy measure for a system disturbance. A bivariate correlation was performed between *robber angle* and *distance* that included all subjects. As predicted, the correlation was not significant ($r_{241} = -0.059, p = 0.360$)(see Figure 2.2 A), even though there was sufficient power to detect even a small effect ($P(r \geq 0.2) = 0.933$).

Second, it was expected that the gross behaviour of the robber should be correlated with the behaviour of the dodger – that is, a disturbance to the system was expected to result in compensatory action by the dodger. Therefore, *robber angle* was correlated with *dodger angle*. Again, as predicted, the correlation was significant ($r_{247} = 0.449, p < 0.001$)(see Figure 2.2B).

To verify that the pattern of results reported above held for every animal and was not an aggregate artefact, *robber angle* was correlated with *distance* and *robber angle* with *dodger angle* for each of the rats individually. As was the case for the pooled data, each subject had a stronger correlation between *robber angle* and *dodger angle* than between *distance* and *robber angle* (see Table 2.2).

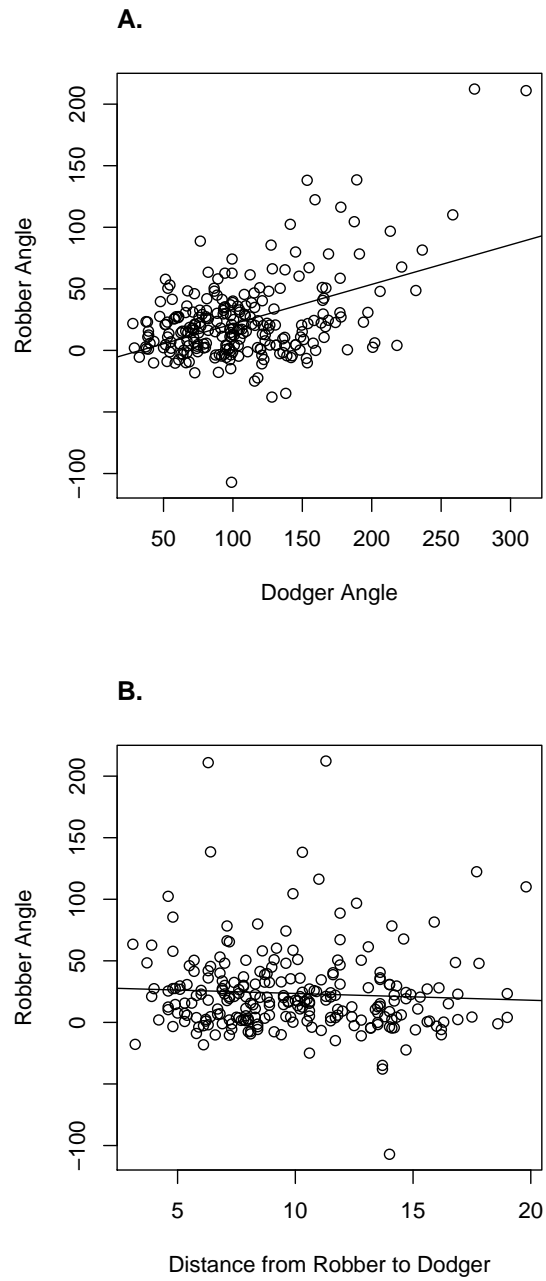


Figure 2.2: Correlations between a) robber angle and dodger angle, and b) robber angle and inter-animal distance.

Table 2.2: Correlations for individual animals.

Subject	Angles ¹	Angle vs. Distance ²
1	0.573	0.095
2	0.717	-0.295
3	0.476	-0.047
4	0.422	-0.160
5	0.363	0.022
6	0.258	0.039
7	0.219	-0.205

¹ Robber angle vs. dodger angle.

² Robber angle vs. inter-animal distance.

Variability

Because *food type* and *partner identity* might not be the only two sources of contextual information that rats use in order to calculate dodge angle, it needed to be determined whether or not *dodger angle* was more variable than *distance*. To test the prediction that *dodger angle* is more variable than *distance*, coefficients of variation were calculated for *distance* and for *dodger angle* for each animal. The coefficient of variation, $\frac{\sigma}{\mu}$, yields a unitless index of variation that can be used to compare distributions measured with different scales (Lehner, 1996). A paired-samples *t*-test was run, and the results indicated that *dodger angle* was more variable than *distance*, $t_6 = 7.849, p < 0.001$. Once again, the pattern observed for each individual animal supports the conclusion based on the pooled data that *dodger angle* was indeed more variable than *distance* (see Table 2.3).

Individual Differences and Contextual Information

An ANOVA with initial inter-animal distance, d_i , as the dependent measure, and *rat ID*, *partner* and *food type* as independent measures, revealed a significant effect of *rat ID* ($F_{6,221} = 9.26, p < 0.001$) (see Figure 2.3). This suggests that there are individual differences with respect to the inter-animal distance when evasive dodges are initiated. *Food type* also produced a significant main effect ($F_{1,6} = 57.39, p < 0.001$), indicating that the inter-animal distance when dodges were initiated varied as a function of the type of food being consumed (see Figure 2.4). The *rat ID* by *partner* by *food type* interaction was also significant ($F_{6,221} = 2.41, p = 0.03$), suggesting a more complex interplay between individual differences and contextual information in the “setting” of the distance at which dodgers began dodging.

An ANOVA, with final inter-animal distance, d_f , as the dependent measure, and

Table 2.3: Coefficients of Variation

Subject	Distance	Dodger Angle
1	0.274	0.374
2	0.320	0.472
3	0.244	0.391
4	0.278	0.345
5	0.346	0.439
6	0.339	0.401
7	0.353	0.467
Mean	0.308	0.413

rat ID, *partner* and *food type* as independent measures revealed that *rat ID* produced the only significant main effect ($F_{6,215} = 17.47, p < 0.001$), suggesting that dodgers differ with respect to their “preferred” inter-animal distance – the distance at which the subjects “felt comfortable” enough to stop dodging (see Figure 2.3). The *rat ID* by *partner* by *food type* interaction was also significant ($F_{6,215} = 2.77, p = 0.013$), indicating that “preferred” inter-animal distance can be modified by contextual information, but that the modification does not occur in a straightforward manner.

Because what was of interest here was whether or not dodgers used contextual information to make dodging decisions, and whether or not dodgers showed individual differences in dodge characteristics, no analyses were run using d_r as a dependent measure. This was because d_r , the distance between the robber and dodger when the robber stopped robbing, reflects behavioural decisions of the robber, and is likely only minimally influenced by any dodging characteristics (the metric of interest) (Figure 2.3).

In order to test further that dodgers were paying more attention to robber movement than to any contextual information, partial and semi-partial correlations for d_i , d_r , and d_f were performed on each subject, and on all of the subjects as a group (see Table 2.4). As expected, the overall partial correlation between d_i and d_f , with the correlations with d_r removed from both variables, was not significant ($r_{241} = 0.112, p > 0.05$), indicating that, once the influence of d_r was removed from both variables, d_f was not dependent on d_i . In other words, the “preferred” inter-animal distance attained by the dodger was not related to the initial inter-animal distance at which dodges were initiated. However, the partial correlation between d_f and d_r , having removed the correlation between d_i from both variables, was significant ($r_{241} = 0.357, p < 0.0001, r^2 = 0.128$), indicating that d_f *does* depend on d_r . This overall pattern of results was mirrored in the partial correlations performed on each

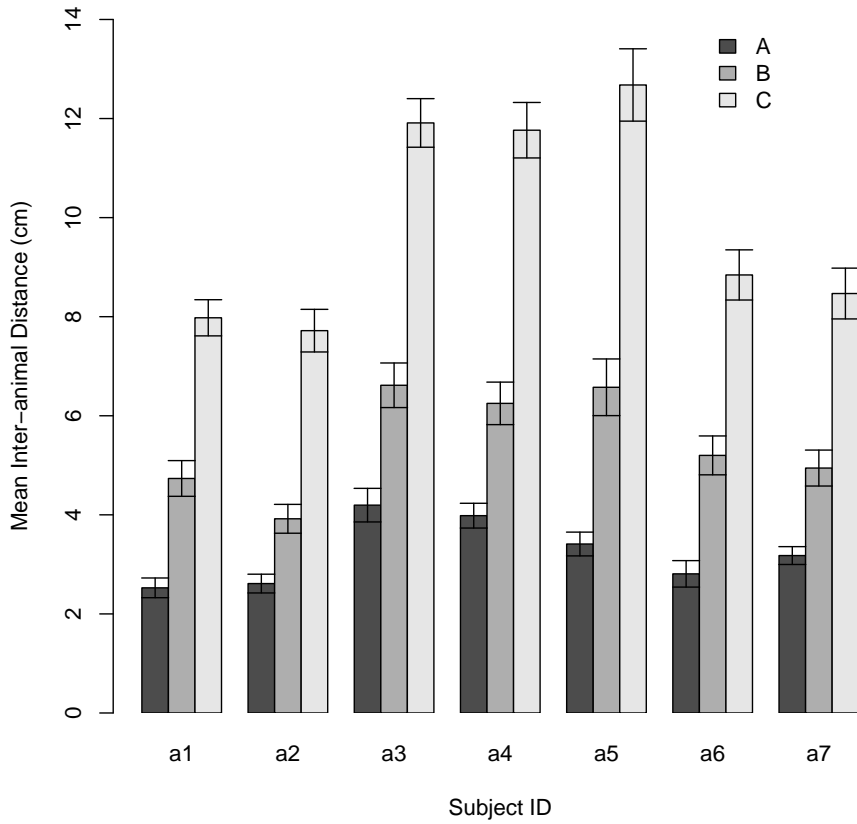


Figure 2.3: Mean inter-animal distance by subject at three timepoints: A) initial inter-animal distance, d_i ; B) inter-animal distance when robber stopped robbing, d_r ; C) inter-animal distance when dodger stopped dodging, d_f . Error bars are 95% confidence intervals.

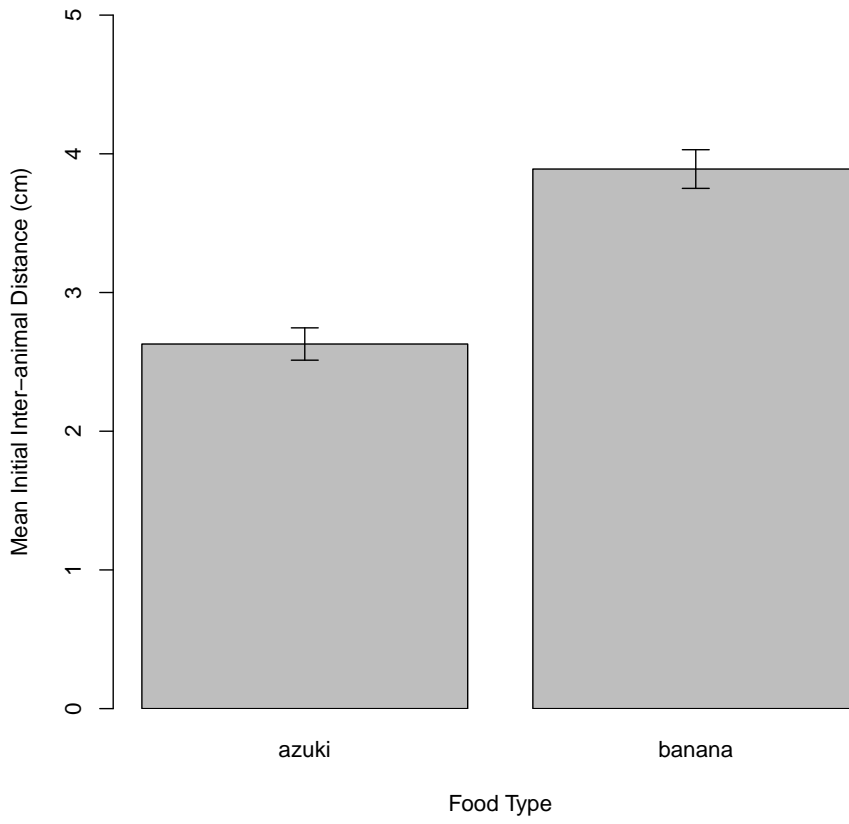


Figure 2.4: Initial inter-animal distance, d_i , as a function of food type. Error bars are 95% confidence intervals.

animal individually (see Table 2.4).

The semi-partial correlation between d_i and d_f , with the correlation between d_r and d_f removed, was significant ($r_{241} = 0.138, p = 0.032$); however, $r^2 = 0.019$, indicating that d_i accounts for less than 2% of the variation in d_f once the influence of d_r on d_f was removed. Finally, the semi-partial correlation between d_f and d_r , with the correlation between d_i and d_f removed, was significant ($r_{241} = 0.462, p < 0.0001, r^2 = 0.213$), further supporting the hypothesis that d_f is dependent on d_r , but not on d_i . As with the partial correlations, the semi-partial correlations that were performed on each subject separately showed the same pattern as the overall semi-partial correlations (see Table 2.4).

Discussion

This experiment showed, by demonstrating the lack of a correlation between inter-animal distance and system disturbances (the movement of the robber), that inter-animal distance was being controlled by the dodger. Additionally, because the movement of the robber and the movement of the dodger *were* correlated with one another, the lack of a correlation between distance and robber movement is less likely to have simply been the result of experimental error.

Further supporting the interpretation of inter-animal distance as a controlled variable, the variability in inter-animal distance was significantly less than the variability in dodge angle, and this pattern was seen for every subject. If the rats really had been calculating fixed dodging angles, then the distance between the two animals should have been irrelevant, and hence more variable than the angles of the dodges traversed by the dodger.

It stands to reason that if a rule such as, “maintain a minimum distance,” is truly

Table 2.4: Partial and Semi-Partial Correlations

Subject	$r_{(d_f d_i).d_r}$ ¹	$r_{(d_f d_i).d_r}$ ²	$r_{d_i(d_f.d_r)}$ ³	$r_{d_r(d_f.d_i)}$ ⁴
1	-0.114	0.484	-0.183	0.571
2	0.127	0.284	0.138	0.336
3	0.138	0.299	0.155	0.363
4	-0.069	0.424	-0.085	0.480
5	0.187	0.340	0.228	0.447
6	0.030	0.218	0.029	0.218
7	0.023	0.140	0.020	0.132
Overall	0.112	0.357	0.138	0.462

¹ Partial correlation between d_f and d_i , with each of their correlations with d_r removed.

² Partial correlation between d_f and d_r , with each of their correlations with d_i removed.

³ Semi-partial correlation between d_f and d_i , with the correlation between d_f and d_r removed.

⁴ Semi-partial correlation between d_f and d_r , with the correlation between d_f and d_i removed.

what is being used, that there necessarily exists some specific reference value for “minimum distance” within each animal, and that the specific value of this reference level need not be the same across individuals. This was clear both for the initial distance between the robber and dodger (i.e., when the dodger started dodging) and for the final distance achieved (i.e., the distance between the animals when the dodge was completed) (Figure 2.3).

The current experiment also shows that previously proposed notions regarding how contextual information influences dodging are likely incomplete. That is, it had been suggested that individual animals use contextual information *only* – type of food, identity of partner, etc. – to make appropriate decisions about how to evade robbers. Although the rats used *some* contextual information when making dodging decisions (i.e., type of food influenced the inter-animal distance at which dodges were initiated), they did not use it in the straightforward manner that had been suggested (i.e., to directly calculate a fixed dodge angle). Instead, contextual information was used to modify the internal reference value of the controlled variable, inter-animal distance. Partner size was also used, in an even less straightforward manner, as contextual information that influenced the reference value for inter-animal distance, and this was seen in the significant *partner* by *food* by *rat ID* interaction.

The specific movements of the robber were also found to influence the subject’s internal reference value for inter-animal distance, as was demonstrated by the significant partial and semi-partial correlations between d_f and d_r , after having taken into account the influence of d_i on one or both variables. It was demonstrated earlier that “static” contextual information (contextual information that does not change throughout the dodge) – in this case, *food type* – influenced the reference value for inter-animal distance at the beginning of the dodge. However, it appears that, as the dodge progressed, what the robber was doing had an increasing impact on the

behaviour of the dodger. This makes sense in that, in the absence of any other information, static contextual information provides the best *a priori* estimation of behaviour that will be effective. A truly effective evasion, however, would incorporate into the strategy additional information (the specific behaviour of the robber) as it becomes available.

These results demonstrate that a simple mechanism underlying what initially appears to be a complex behaviour that accounts for all of the behaviour present, including phenomena that might previously have been attributed to foresight, planning, and computation. That is not to say that computation does not take place – just that computation based on abstract symbol manipulation is not required (i.e., rats do not need to calculate fixed dodging angles based on contextual cues). Instead, the computations performed are organic, dynamic and embedded in the system. Instead of relying on ever-more-complicated cognitive constructions, efficient, effective and seemingly complex behaviour can be explained by appealing to arguably more parsimonious simple rules. And given the breadth of examples of apparently similar phenomena in a variety of organisms, it is likely that the cybernetic approach can be applied to a wide range of behaviour heretofore thought to be the result of complicated cognitive computations.

Chapter 3

HOW IS A CRICKET LIKE A RAT? INSIGHTS FROM THE APPLICATION OF CYBERNETICS TO EVASIVE FOOD PROTECTIVE BEHAVIOUR

In the previous chapter, it was demonstrated that a simple cybernetic rule, gaining and maintaining a preferred inter-animal distance, can account for much of the variability in dodging by rats. In this chapter, the field cricket, *Teleogryllus oceanicus* was used to test whether or not the same or similar cybernetic, rather than S-R rules are employed by animals of different lineages and body plans by testing whether or not crickets operate using the same principles in similar contexts. Field crickets were chosen to conduct this investigation for several reasons. First, cybernetic rules have previously been used to explain some aspects of cricket behaviour (Pellis et al., 2009). Second, as crickets have also been observed to perform dodges in response to robbing attempts made by other crickets (unpublished observations), the same overt behaviour could be compared across the two species. Third, although they seem to perform dodges to protect food in a similar fashion, crickets have a different body plan than rats, with Arthropoda and Chordata having diverged at least 500 million years ago (Ayala & Rzhetsky, 1998). For example, rats use their forepaws to hold their food, requiring them not only to shift their body weight onto their hind legs, but also limiting their defensive movements to only two of their feet (Whishaw, 1988). In contrast, crickets hold their food with their mandibles (Manton & Harding, 1964), leaving all six limbs free for locomotion when evading the robber. Such differences in posture and motor strategies across the species could produce significant differences in the organization of food protection.

Four hypotheses were tested in the present chapter. The first two major hypotheses were: H1) that crickets use a cybernetic rule during food protective behaviour,

and H2) that crickets employ the same “gain and maintain a minimum inter-animal distance” rule in the same way that it is used by rats – that is, despite all of the differences between the two species, crickets and rats defend food using the same underlying principle (see Table 3.1). The second two minor hypotheses, dependent on the main hypotheses, were: H3) that, in addition to controlling inter-animal distance, defenders also control their body orientation relative to robbers, and H4) that approach orientation of the robber is related to the type of evasive action employed by defenders (see Table 3.1).

One issue with respect to H1 is that for crickets, and for arthropods in general, rapid defensive and offensive actions are thought to be ballistic, allowing little opportunity for modification once their behaviour is initiated (Schöne, H., 1984; W. J. Bell, 1991). If arthropods act solely in a ballistic fashion, cybernetic mechanisms, because they rely on feedback, could not be a fundamental component of the organization of cricket behaviour. However, the same S-R, non-updateable type of rule was also previously thought to apply to the dodging behaviour of rats (Whishaw & Gorny, 1994). Fortunately, ballistic and cybernetic rules produce distinct behavioural outcomes. If defender crickets respond only ballistically to the stimulus of the robber’s approach, their behaviour should be relatively stable through time and independent of the specific movements used by the robber. Conversely, if crickets are able to use feedback to compensate for changing conditions, then defender responses should not only be variable, but should also be correlated with the specific movements used by the robber.

With respect to H2, if the goal of the defender is to control a 2-dimensional inter-animal distance, then, as was discovered in the rats, inter-animal distance should be less variable than any other measured variable (H. C. Bell & Pellis, 2011). Further, inter-animal distance should not be correlated with the actions of the robber – that is,

inter-animal distance should remain relatively constant even though the movements of the robber do not – which was also characteristic of the rat dodging system (H. C. Bell & Pellis, 2011). Finally, the same pattern of relationships between all measured variables that was present in the rats, such as a correlation between the movement of the robber and the movement of the defender, should also be present in the crickets.

With respect to H3, our initial observations indicated that, like rats (Field et al., 1996), cricket defenders faced away from robbers at the end of interactions. It was predicted that, if defender body orientation relative to robber were controlled – specifically, if the “facing away” orientation were preferred – then the “facing away” orientation should be favoured over other defender orientations at the end of interactions.

With respect to H4, unlike rats, crickets do not rely solely on dodging (i.e., pivoting away from the robber) as an evasive strategy. In many instances, crickets simply walk or run directly away from the robber. Crickets also differ from rats in the variety of approach angles that are used by robbers. A rat primarily approaches another rat holding food from a parallel orientation, with its head moving along the flank, in the direction of the defender’s head. Thus, in the majority of encounters, the defender is confronted with a similar robbing orientation. In contrast, as crickets in their approach to rob do so from virtually any orientation, this difference could account for the differing evasive strategies employed by crickets. Therefore, it was predicted that the point on the body of the defender targeted by the robber is related to the type of evasive strategy adopted (i.e., dodging versus running).

Method

Subjects

Ninety-two female field crickets (46 pairs) (*Teleogryllus oceanicus*), born and raised at the University of Lethbridge, were used in total; however, not all animals were used for every analysis. All animals were housed in rooms maintained at 25 °C and 70% relative humidity, on a 12h light/dark schedule, with lights on at 12:00h. Hatchlings were placed into large plastic bins (51cm long x 37.5cm wide x 35.5cm high) containing layers of cardboard egg cartons for shelter, glass shell vials filled with water and stoppered with cotton for moisture, and ground and pelleted cat chow (Iams® Original with Chicken) *ad libitum*. Water vials were changed weekly and additional food was given as needed. Female nymphs were separated from the colony during their penultimate instar (final juvenile instar) and were then housed in groups of 10 to 20 individuals in containers (29.5cm long x 19cm wide x 12.5cm high) with layers of cardboard egg cartons for shelter, with water and food provided as before. Following their final moult, the animals were placed, individually, into round plastic containers (9cm diameter x 8cm high). Half of the animals were randomly assigned to be subjects, defending the food, and the other half were assigned to be robbers. Half of both the defenders and the robbers were marked with a dab of typewriter correction fluid on their pronotum. Each cricket destined to be a defender was randomly assigned a partner that would act as the robber. Although isolated, the animals were housed in the colony room from which they would have been able to see, hear, and smell conspecifics. Thirteen of the original 46 pairs were eventually excluded from the study due either to one of the pair dying or escaping, or to lack of sufficient data to include in the analyses (i.e., too few interactions).

Given that different cricket pairs interacted at different rates, the number of in-

teractions contributed by each pair differed. To ensure that the patterns discerned were not a product of over-contribution by a few pairs, only four examples from each pair were used, and so, all pairs contributed an equal number of interactions, yielding a balanced data set for testing. Moreover, since some pairs interacted more than four times, to avoid bias in selecting examples, for each pair, the first four interactions for each type of defence that met the criteria of occurring away from the walls were selected. Comparisons involving the whole data set (unbalanced) are only presented if validated by a balanced analysis.

Testing

Robbers and defenders were deprived of food for six days prior to testing, but given water *ad libitum* during that time. Prior to conducting this experiment, female field crickets were subjected to varying degrees of food deprivation. Preliminary testing showed that six days of food deprivation was required to produce sufficient motivation for animals to attempt to steal food from one another without affecting the mortality rate (unpublished observations). Previous research also indicated that field crickets can be deprived of food for up to 7 days without an increase in mortality (Adamo & Hoy, 1995). On the testing day, each subject (the defender) and its partner (the robber) were introduced into the testing chamber – a 12.6cm diameter Plexiglas® cylinder situated on top of a clear platform with a mirror underneath at a 45° angle. The animals were allowed to habituate to the enclosure for two minutes, at which point, a small piece of cat chow (approximately 3mm in diameter) was placed, using forceps, in front of the defender. The trial began when the defender picked up the food with its mandibles, and ended when the defender had finished eating the food, when the robber succeeded in stealing the food from the defender, or when 12 minutes

had elapsed. Trials were conducted in the dark, and were filmed from below through the mirror using a Sony High Definition digital video camera in the infrared spectrum. Following the trials, the animals were returned to their home containers, in which cat chow and water were available *ad libitum* for 1 day, after which the cat chow was removed, and the animals were again deprived of food for 6 days. Trials were always performed using the same pairings of animals, and the cycle of trials was repeated 3 times. Following the experiment, approximately half of the surviving crickets were returned to the main colony, and the other half were euthanized by freezing and placed in individual vials containing a 70% EtOH solution so that body size measurements could be taken. The crickets that were euthanized were chosen randomly from the entire sample. Both members of the original pair were euthanized so that there was an equal number of robbers and defenders, and so that body size differences within pairs could be compared. As the body size measurements were used for a subsequent study on combat behaviour, they are not reported here.

Video Analysis

All video was initially viewed at full speed. Instances of fighting and evading were counted (see below for descriptions of the different types of behaviours in these two categories). All instances of evading were digitally clipped from the main video, and individually analysed frame-by-frame using tracking software (Vicon Motus, Vicon Motion Systems, Colorado, USA). Several measurements were made with that system (see below). Although the majority of the analysis was performed by the author, a subset of the interactions were re-analysed by a naive observer and scores were correlated to ensure that the full analysis was unbiased. Specifically, the angles swept through by both the robber and dodger were correlated across observers during

dodging interactions ($r_4 = .871, p = .02$ and $r_4 = .862, p = .03$ respectively), indicating that when the animals were digitally tracked, the starting and end points of the tracking were the same across observers. Paired Wilcoxon tests showed that the mean angles traversed by both robbers and dodgers did not differ between observers (both results were $V = 5, p = 0.3125$), which was further supported by dependent-samples t-tests ($t_5 = -1.609, p = .169$ for dodge angle and $t_5 = -1.403, p = .216$ for robbing angle). Unfortunately, because the Motus system was calibrated differently, distance measurements could not be directly compared across observers. To ensure that the main analysis was consistent, the balanced subset of dodges was split into two groups. Welch's independent samples t-tests showed no difference in dodging or robbing angles ($t_{52.655} = -.5307, p = 0.5979$ and $t_{46.317} = 0.2509, p = 0.803$ respectively). Additionally, no difference was seen across samples in the mean inter-animal distance ($t_{53.944} = 1.229, p = 0.2245$).

Differing Types of Defence

Crickets use two main types of defensive techniques during food protection, fighting and evading. Fighting involves one of two tactics, grappling, in which the defender faces the robber and the pair wrestle with interlocked mouthparts, and kicking, in which the defender faces away from the robber and kicks at it with its hind legs (Adamo & Hoy, 1995). Evading involves the defender either laterally swerving away from the robber (i.e., dodging) or running directly away from the robber. The number of different types of defences that were used by each animal was counted. Because what was of interest was comparing cricket to rat behaviour, and because rats use only evasion to protect food items, fighting in crickets will not be discussed further.

Dodging and Running

When comparing species, a methodological problem arises. Species may differ in the behaviour of interest because the opportunity to perform that behaviour may differ between species. For example, in fighting in rodents, the actual frequency of use of class-common defensive tactics may vary, in part, because of the location in which the attacks occur. If a defender remains wedged in a corner facing its opponent, the kind of attack possible differs from a case in which the defender flees in an open space. Similarly, the tactic employed may differ in structure if used in a constrained versus an unconstrained context (Pellis, Pellis, Pierce, & Dewsbury, 1992). In the robbing and dodging paradigm, these problems can be solved by only comparing dodges during which rats are at least one body length away from a wall (Field et al., 1996). This was the methodological approach used for the previous study of robbing and dodging in rats (H. C. Bell & Pellis, 2011), and is the one that was used for the crickets in the present study.

The one difference between rats and crickets, as already noted, was that in crickets, attackers can approach from any orientation, rather than the more restricted one typical of rats. To adjust for this difference, in crickets, evasions from different orientations were separately measured and evaluated in case the orientation of approach produced differences in evasive strategies. Once the evasions had been selected, Vicon Motus software was used to track the position, digitally, of the most anterior point on the head and the most posterior point on the abdomen of both the robber and defender during the behaviour. The evasion was said to have begun in the first frame in which the defender began to move. The evasion was completed in the final frame of the defender's movement. As is the case in rats, the robber always stopped robbing before the defender stopped dodging or running.

Several measures were extracted from the digitized tracking data. For dodging only, the angle traversed by the defender (the dodging angle) was used, but unlike in our study on rats, the angle traversed by the robber was not used. This was because, in rats, the robber tends to pivot along with the defender, whereas in crickets, the robber tends either to stop once the defender begins to pivot, or follows in a straight line. Therefore, measuring the angle traversed by the robber would not have been meaningful. For all evasions, be they dodges or runs, the path lengths and maximum instantaneous velocities of both the robber and the defender were measured, as well as the distance between the heads of the two animals (the inter-animal distance). In order to ensure that the data on the crickets could be compared to that of the rats, the sequences of robbing and dodging from H. C. Bell and Pellis (2011) were re-analysed to measure the path length traversed by both robbers and dodgers.

As in the previous chapter, inter-animal distance was used both as a continuous measure (tracked through the entirety of the dodge) and at three specific time points. The initial inter-animal distance, d_i , was measured in the first frame in which the evasion began. The inter-animal distance when the robber stopped robbing, d_r , was measured in the last frame in which the robber pursued the defender. The final inter-animal distance, d_f , was measured in the last frame in which the defender evaded the robber, and was taken as the measurement of the distance between the current position of the defender, and the position of the robber at the d_r time point. d_f was calculated in this way to account for any lag time in the updating of the information available to the defender. In addition, we wanted to ensure that any irrelevant movement on the part of the robber after the end of its pursuit of the defender, such as walking in the opposite direction, would not be included in the measurement of d_f . See Chapter 2 for further details.

In addition to the tracking data, the orientation of the robber with respect to

the defender was recorded both at the initiation of the evasive manoeuvres and at their completion. The initiation and the termination of the evasive manoeuvre were defined, respectively, as the first frame in which the defender began to move, and the last frame in which defender moved (see above). Three categories of approach orientation were also recorded: the robber could approach from the front (f), in which the robber was facing the head of the defender, the side (s), in which the robber was facing the defender's flank, or from the rear (r), in which the robber was facing the rear of the defender's abdomen (see Figure 3.1).

Results

The results are presented below. See Table 3.1 for a summary of the predictions made and the results obtained.

Are Crickets Cybernetic Agents?

If a ballistic rather than a cybernetic rule were being used by the crickets, the amount of movement made by each defender in response to a robbing attempt should be static. However, for both running and dodging, in balanced subsets of the first four interactions per animal (the same subsets used hereafter, unless otherwise indicated), the amount of movement (path length) was variable (see Figure 3.2). Overall, running crickets had path lengths ranging from 0.919 to 31.188 cm, with a mean of 10.008 cm and a standard deviation of 4.685 cm. The coefficient of variation, $\frac{\sigma}{\mu}$, for running crickets was 0.465. For dodging crickets, the range was 3.346 to 28.404 cm, with a mean of 10.424 cm and a standard deviation of 4.708 cm. The coefficient of variation was 0.452.

If, as in rats, dodging were a coupled system in crickets, the movement of the

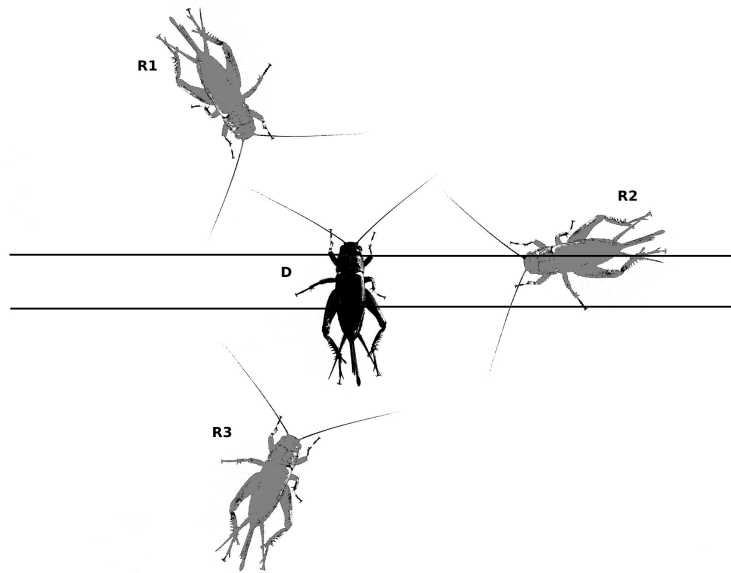


Figure 3.1: Illustration of how the robber orientation was determined. D is the defender. R1 is a robber in the front orientation; R2 is a robber in the side orientation; R3 is a robber in the rear orientation. The horizontal lines denote the cut-off criteria used.

Table 3.1: Competing Hypotheses and Predictions Tested

Hypotheses 1	Crickets use cybernetic rules to protect food from robbers	Crickets use ballistic rules to protect food from robbers
Predictions ¹	<ol style="list-style-type: none"> <i>Defender responses to robbing attempts are variable.</i> <i>Defender and robber movements correlated.</i> 	<ol style="list-style-type: none"> Defender responses to robbing attempts are invariant. Defender and robber movements not correlated.
Hypotheses 2	Crickets use the same rule as rats to protect food from robbers	Crickets do not use the same rule as rats to protect food from robbers
Predictions	<ol style="list-style-type: none"> <i>Inter-animal distance is invariant.</i> <i>Inter-animal distance is not correlated with robber movement.</i> <i>The same pattern of relationships among all variables seen in rats is seen in crickets.</i> 	<ol style="list-style-type: none"> Inter-animal distance is variable. Inter-animal distance is correlated with robber movement. The same pattern of relationships among all variables seen in rats is not seen in crickets.
Hypotheses 3	Defender body orientation is a controlled variable	Defender body orientation is not a controlled variable
Predictions	<ol style="list-style-type: none"> <i>Defender body orientation is the same at the end of every interaction.</i> 	<ol style="list-style-type: none"> Defender body orientation is variable at the end of every interaction.
Hypotheses 4	Initial robber orientation predicts evasive strategy used	Initial robber orientation has no bearing on evasive strategy used
Predictions	<ol style="list-style-type: none"> <i>Initial robber orientation and evasive strategy are correlated.</i> 	<ol style="list-style-type: none"> Initial robber orientation and evasive strategy are not correlated.

¹ Predictions that were supported are *italicized*.

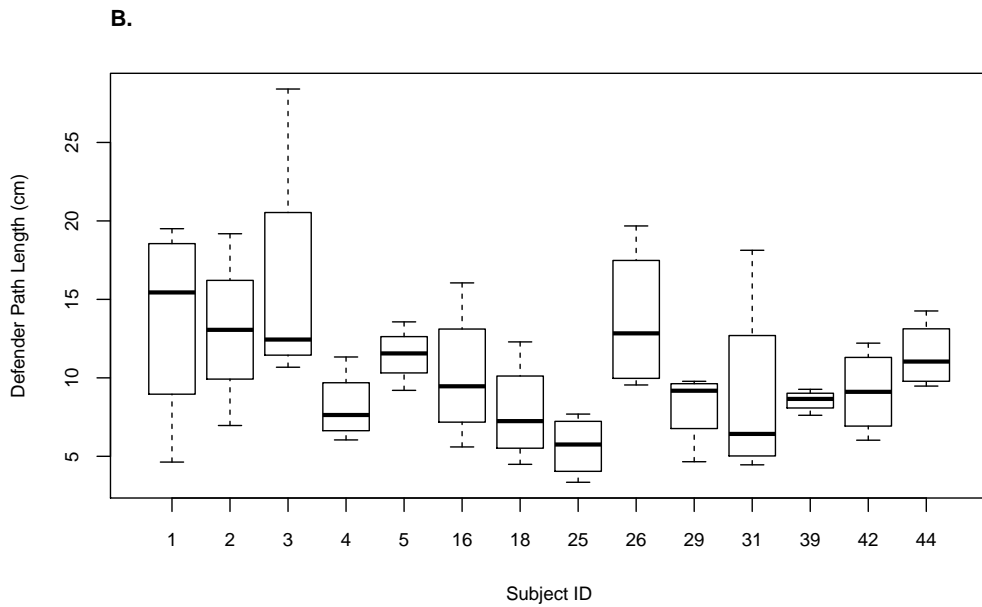
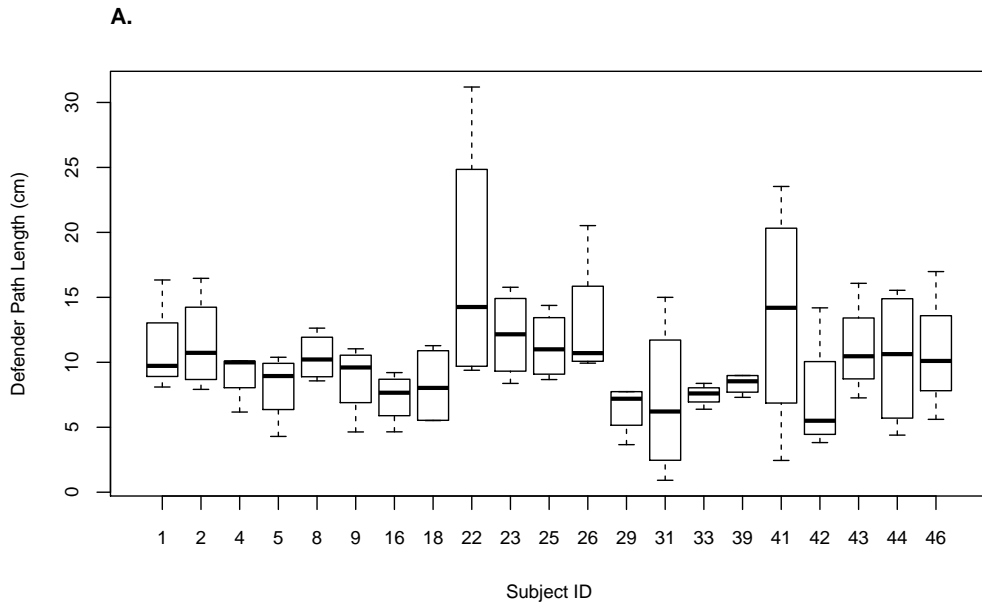


Figure 3.2: A. Variability in path length in running crickets over 4 interactions. B. Variability in path length in dodging crickets over 4 interactions. Whiskers indicate range, boxes are first and third quartiles, and centre bar is the median.

defender and the movement of the robber should be correlated. Using the path length of each animal as a measure of movement, the prediction was supported for the whole dataset (running plus dodging) ($r_{279} = .634, p < .001$) (see Figure 3.3A.). For dodging alone, the rob path vs. defender path analysis yielded $r_{54} = .548, p < .001$, and for running, the rob path vs. defender path produced $r_{82} = .732, p < .001$.

Are Crickets the Same as Rats?

It was predicted that, if, as was the case in rats, inter-animal distance were a controlled variable, the variation for inter-animal distance should be less than for any other measure. The coefficient of variation, $\frac{\sigma}{\mu}$, was used, so that variables measured using different scales could be directly compared (Lehner, 1996). The prediction was supported (Table 3.2).

Also, as in rats, if inter-animal distance were a controlled variable, it should not be correlated with the movement of the robber. As predicted, the path length of the robber was not correlated with inter-animal distance for the entire data set ($r_{279} = 0.18, p = .762$) (see Figure 3.3B.). The absence of a correlation was not an artefact of having insufficient data, as further analysis showed that there was sufficient power to detect even a small effect ($P(r \geq 0.2) = 0.923$). Additionally, for dodging alone, rob path vs. distance produced $r_{54} = .001, p = .995$, and for running alone, $r_{82} = -.045, p = .681$.

If crickets use the same rule for organizing evasive behaviour as rats, one would expect that the pattern of correlations between certain variables would be the same for crickets as they are for rats. Using path lengths from the rat data (H. C. Bell & Pellis, 2011), the variables measured were compared with the analyses of the crickets when they both dodged and ran. Irrespective of the evasive tactic used, the same

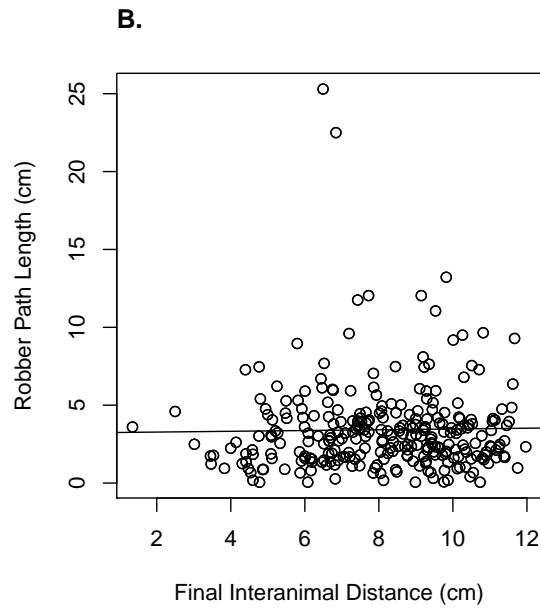
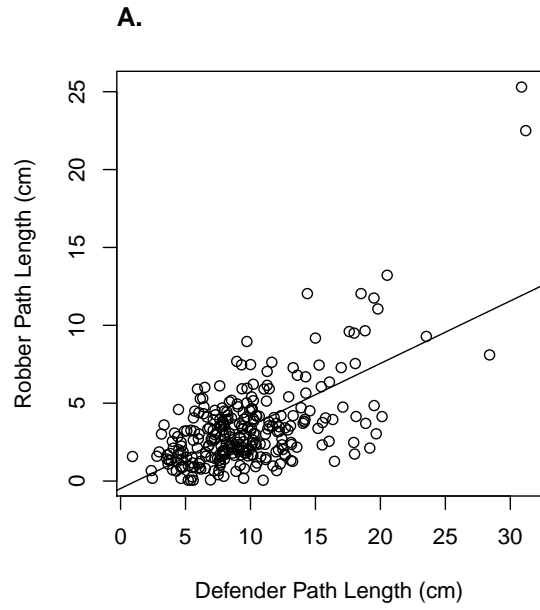


Figure 3.3: A. Correlation between robber path length and defender path length. B. Correlation between robber path length and inter-animal distance.

Table 3.2: Coefficients of Variation, $\frac{\sigma}{\mu}$

Measure	Dodging Crickets ¹	Running Crickets ²	Rats ³
Dodge Path Length (cm)	0.452	0.465	0.463
Rob Path Length (cm)	0.759	0.872	0.867
Distance ⁴ (cm)	0.341	0.193	0.392

¹ $n = 56$

² $n = 84$

³ $n = 246$

⁴ Final inter-animal distance, d_f

pattern was found in crickets, and this pattern matched that of rats when they dodged in almost all comparisons (Table 3.3). The one exception was that, for rats, inter-animal distance and maximum instantaneous dodge velocity were correlated, whereas this was not the case in crickets for either running or dodging.

As with the rats, in the crickets, the initial inter-animal distance, d_i , was shorter than the distance after the robber had ceased robbing, d_r , and these were both shorter than the final inter-animal distance, d_f (see Figure 3.4). The relationships between d_i with d_f , for both the running and dodging conditions, were tested using correlations, neither of which were significant ($r_{82} = .023, p = .832$ and $r_{54} = .094, p = .491$, respectively). This mirrors what was found in rats. In contrast, d_r and d_f were significantly correlated for both running and dodging ($r_{82} = .627, p < .001$ and $r_{54} = .648, p < .001$, respectively), again, as was found in rats.

Is Defender Orientation a Controlled Variable?

It was also predicted that, in addition to inter-animal distance, the final orientation of the defender to the robber would be facing away, as was the case in rats (Field et al., 1996). This is what was found (Table 3.4).

Does Approach Orientation Determine the Type of Evasive Strategy Used?

Using the full data set, the frequencies of the initial orientation of the robber with respect to the defender were recorded (Table 3.5). A Chi-square test of independence indicated that defenders that were approached from the front would almost always use dodging as an evasive strategy; whereas defenders that were approached from the rear would almost always use running as an evasive strategy. Defenders that were

Table 3.3: Pattern of Correlations Among Variables for Crickets and Rats

Measure	Dodging Crickets ¹	Running Crickets ²	Rats ³
Dodge Path x Rob Path	0.548*	0.732*	0.512*
Dodge Velocity ⁴ x Rob Velocity	0.503*	0.321*	0.162*
Dodge Velocity x Dodge Path	0.418*	0.336*	0.451*
Rob Velocity x Rob Path	0.806*	0.635*	0.433*
Distance ⁵ x Dodge Path	0.701*	0.488*	0.540*
Distance x Dodge Velocity	0.244	0.113	0.481*
Distance x Rob Path	0.008	-0.045	-0.083
Distance x Rob Velocity	0.053	0.019	0.008

¹ $n = 56$

² $n = 84$

³ $n = 246$

⁴ Maximum instantaneous velocity

⁵ Inter-animal distance

* Denotes result is significant at $p < .05$

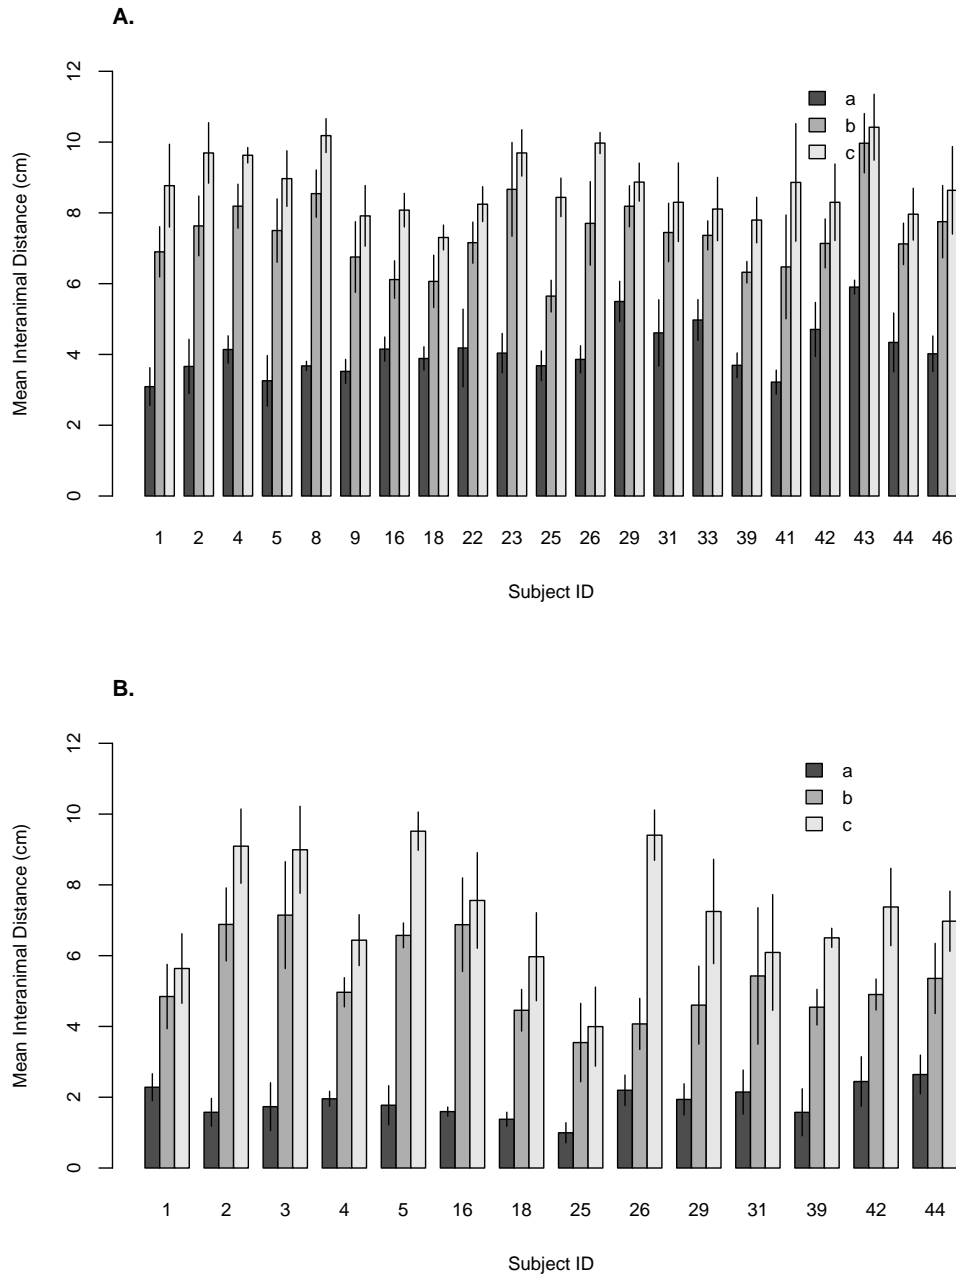


Figure 3.4: A. Mean inter-animal distance for running crickets at three time points: a) d_i ; b) d_r ; and c) d_f . B. Mean inter-animal distance for the same three time points as in A. for dodging crickets.

Table 3.4: Counts of Final Defender Orientations with respect to Robbers

Behaviour	Facing	Side	Away	Total
Run	0	0	192	192
Dodge	0	1	89	90
Total	0	1	281	282

approached from the side were equally likely to dodge or run ($\chi_2^2 = 222.41, p < .001$).

Given that the robbers could initiate attacks from different orientations (Figure 3.1), and that different approach angles were associated with different evasive tactics (Table 3.5), approach angle could potentially influence the ability of the defender to coordinate its movements with those of the robber. As shown in the above analysis, the variation in the distance gained from d_i to d_f is partly dependent on the movement performed by the defender concurrently with that of the robber (d_r). Therefore, the strength of the correlated movements by robbers and defenders should be greatest when the robber approaches from an orientation in which the two animals are equally able to counter each others' movements. Approaching from the side would seem to be such an orientation as the defender is equally likely to run or dodge, but, in either case, as the defender moves, the robber can countermove, keeping a constant orientation. Conversely, frontal approaches invariably lead to dodging, and, as pivoting around a vertical axis moves the head away faster than the walking or running needed to maintain the orientation of the robber, this approach orientation should afford the robber the least opportunity to counter the defender's movements finely. The data support these expectations (Table 3.6). The strongest correlation for defender path versus robber path was for side approaches and the weakest for frontal approaches, with rear approaches yielding an intermediate value. With regard to final distance and movement by the defender (defender path vs. distance), the strongest correlation was when the robber approached from the front. That is, when approached from the front, most of the change in distance from d_i to d_f was determined by the movement of the defender, with little contribution by the movement of the robber.

Table 3.5: Counts of Initial Robber Orientations with respect to Defenders

Behaviour	Front	Side	Rear	Total
Run	0	20	172	192
Dodge	67	20	3	90
Total	67	40	175	282

Table 3.6: Correlations by Approach Orientation

Measure	Side ¹	Front ²	Rear ³
Defender Path x Robber Path	0.825*	0.228	0.586*
Defender Path x Distance	0.459*	0.894*	0.545*

* denotes result is significant at $p < .05$

¹ $df = 22$

² $df = 37$

³ $df = 75$

Discussion

Cybernetic versus Ballistic Movements

Two aspects of the present findings showed that crickets do not behave ballistically when evading robbing attempts by other crickets. First, defender responses, as measured by distance traveled (path length) were variable (Figure 3.2). Given that *T. oceanicus* measure roughly 2cm in length, the measured range of defender path lengths – less than 1cm to more than 30 cm – translates into a range of movement from 50% to 1500% relative to body length. Second, although one would expect that the robber’s path length and defender’s path length should not be correlated if the movements were ballistic (i.e., the movement of the robber should not be influenced by the movement of the defender, or *vice versa*, after the initial “stimulus” event), the data show that a strong correlation exists (see Table 3.3). Moreover, whereas rats showed a strong correlation between dodge velocity and inter-animal distance, crickets did not (Table 3.3), suggesting that the speed of the movement had more of an effect on the outcome of the defending rat’s behaviour than it did for defending crickets. If anything, these data indicate that crickets are less, not more ballistic, in their defensive actions than are rats.

Commonalities in the Defensive Behaviour of Crickets and Rats

As was the case in rats, coefficients of variation, $\frac{\sigma}{\mu}$, indicated that inter-animal distance is less variable than the any other measure (see Table 3.2). Crucially, inter-animal distance was also not correlated with the movement of the robber (Figure 3.3B.) Together, these findings constitute the Test for the Controlled Variable employed in Perceptual Control Theory (PCT), a cybernetic theory of behaviour, which

posits that animals move in order to maintain particular perceptions (i.e., controlled variables) constant (Powers, 1973). The lack of variability relative to all other measured variables indicates that inter-animal distance is controlled by the defender. The lack of a correlation between inter-animal distance and the movement of the robber indicates that, despite disturbances (i.e., robbing attempts), the defender is able to compensate behaviourally in order to defend the state of the inter-animal distance variable. This implies that, rather than responding ballistically to robbing attempts by engaging in stereotypical escape strategies, defenders are constantly monitoring the position of the robber and updating their behaviour accordingly.

The pattern of correlations seen between specific variables in the rat system were the same as those seen in the cricket system, even when the crickets used a different evasive strategy as compared to the rats (Table 3.3). A finding not seemingly consistent with the cybernetic perspective was that defender path length was always strongly correlated with inter-animal distance. Indeed, re-examination of the rat data from H. C. Bell and Pellis (2011), showed the same strong correlation in rats as well as crickets (see Table 3.3). These data indicate that the defender's movement predicts inter-animal distance as much as the robber's movement (see Figure 3.3). That is, the gained and maintained inter-animal distance resulted, in part, from the amount of movement by the defender rather than as movement arising in compensation for the movement by the robber. One way to reconcile these findings is to recognize that the final inter-animal distance is composed of two main components – moving to a preferred inter-animal distance and moving in response to the approaching movement by the robber. This is reflected in the fact that the defender continues to increase inter-animal distance after the robber has ceased its robbing movements (see Figure 3.4). Again, this was also found for rats (H. C. Bell & Pellis, 2011).

The Importance of Orientation

Another possible controlled variable was the final orientation of the robber with respect to the defender. In all but one of the interactions, the defender was facing away from the robber at the end of the encounter, even though encounters were begun in an almost infinite variety of configurations. PCT and cybernetic theories in general, specify that controlled variables can be arranged hierarchically, thus enabling these types of theories to move beyond explaining isolated components of behaviour to explaining, potentially, most aspects of an organism's behaviour (Powers, 1973). Specifically, controlled variables that are lower in the hierarchy must be controlled before higher-order variables can be controlled. Therefore, the finding that the crickets end their dodges with a specific orientation to the opponent could reflect one of two possible mechanisms. (1) The final orientation and the inter-animal distance are controlled simultaneously. That is, they are at the same level on the hierarchy, arranged laterally. (2) The inter-animal distance is at a different level on the perceptual hierarchy than the final orientation. That is, the rule for inter-animal distance and that for final orientation are fulfilled at different times during the evasion. Indeed, it is more likely that distance and orientation are at different levels, given that the "away" orientation is always achieved before the preferred distance. This possibility is supported by the findings in rats, in which the final orientation can be dissociated from the dodging away from the robber (Pellis et al., 1999).

Whereas rats, in most cases, tend to rob from the same approach orientation (Whishaw, 1988), crickets will approach from any orientation (see Table 3.5). Therefore, if the correlation between the defender's movement and inter-animal distance reflects the part of the movement required to gain the preferred distance, then different approach orientations should affect the defender's ability to reach that distance,

and hence, the amount of movement required. When approached from the front, the defender is most likely to dodge, and when approached from the rear, the defender is most likely to run away (Table 3.5). In both cases, the defensive action orients the posterior of the abdomen towards the robber, the preferred end orientation for defenders (see Table 3.4). In addition, these two evasive strategies are also the most efficient means, in each instance, of gaining and maintaining the preferred inter-animal distance.

In contrast, when approached from the side, the defender is equally likely to dodge or run (Table 3.5) to reach the facing away orientation (Table 3.4). Conversely, though, from the side approach, the robber is also better placed to counter the movements of the defender. That being the case, the moves and counter-moves by robbers and defenders should be most strongly correlated when robbers approach from the side (Table 3.6). Indeed, the weakest correlation is when the robber approaches from the front, as a rapid dodge blocks the robber's approach. Although, in order to gain the preferred inter-animal distance, there has to be some degree of correlation between the defender's movement and inter-animal distance, given the counter-moves afforded the robber from the side approach, this correlation should be strongest for side approaches and weakest for frontal approaches, in which the robber has less capacity to counter the defender's moves – and this is what was found (Table 3.6).

To gain and maintain the preferred inter-animal distance, the defender has to move away from the robber sufficiently to gain that distance and then an additional amount to maintain that distance in compensation for the continued approach of the robber. Therefore, both the correlation between the robber's and defender's movements and the correlation between the defender's movement and inter-animal distance are consistent with the defender gaining and maintaining the preferred inter-

animal distance. These movements contribute to maintaining a constant inter-animal distance despite moves and counter-moves by the robber (Figure 3.3). Thus, all the major facets of the cricket's food defence behaviour can be explained in cybernetic terms, with the defender *gaining* and *maintaining* a preferred inter-animal distance. In this regard, crickets and rats organize their food defence behaviour in the same manner.

As was the case in rats, d_i was not correlated with d_f , but d_r was correlated with d_f . This indicates that the behaviour of the defending cricket at the initiation of an evasion is not dependent on the movement of the robber, but by the time the evasion has ended, the defender is using the movement of the robber to make its decision about when to stop moving. As with the rats, it is likely that, in the absence of information from the robber, the dodging strategy is fairly stereotypical, but that as information from the robber becomes available, the defender will use that information to make decisions. That is, given a particular approach orientation, the defender will use a particular evasive strategy initially, but as the robber continues to move, the defender will compensate for those movements.

The alternative to tracking the robber's movements is to use a ballistic action that is either reflexive – simply respond to the looming stimulus – or some calculated response based on contextual cues (e.g., Whishaw & Gorny, 1994). The tests of the predictions support a cybernetic tracking mechanism (see Table 3.1). Second, there could be alternative perceptions to track (see e.g., Cziko, 2000; Marken, 1988, 2009; Powers, 1973), and indeed, it seems that there are at least two – distance and orientation – and these combine in ways to increase the variability of the movement performed in evasions.

It is possible that the behaviour seen in the crickets might actually result from two competing motivations that just happen to balance out in this study. The two

motivations are that for eating the catfood, and that for cannibalizing the other cricket. Given that food-deprived female crickets of some species have been known to readily cannibalize conspecifics of similar size (unpublished observations), it is certainly a possibility. However, we think this is unlikely given that, before the food was introduced to the arena, the pairs of crickets were habituated to the enclosure together. During that time, no attempts were observed by any cricket to cannibalize another in any of the trials, nor were any attempts observed during the pre-testing we did to determine the appropriate level of food deprivation required to produce robbing and dodging behaviour. That is, not a single incidence of cannibalism was observed in more than 60 pairs of crickets (including “pre-test” pairs).

The present results mirror those of another instance in which both mammals and arthropods follow the same simple cybernetic rule. When catching a ball, humans keep the orientation of the ball constant, relative to the horizon, and speed up or slow down their movement accordingly (McBeath et al., 1995). Similarly, when dogs chase and catch frisbees, they use the same rule (Schaffer et al., 2004), as do dragonflies catching prey (Olberg, Worthington, & Venator, 2000). That is, in all these cases, including crickets and rats protecting a food item, the animals keep a simple perception constant and in maintaining this constancy, variable behaviour is produced (Powers, 1973). Although there may be differences in both sensory and motor capabilities that make for subtle differences in being able to detect disturbances and engage compensatory action across species – especially for species spanning different phyla – the present work, along with that cited, strongly supports the idea that many organisms use simple cybernetic principles in structuring at least some of their interactions with the world, a point repeatedly converged upon by both physiologists (e.g., Bernstein, 1967; von Holst & Mittelstaedt, 1950) and ethologists (e.g., Golani, 1976; Uexküll, 1921).

Like rats, crickets are able to protect food from being stolen by other crickets by using evasive strategies. The two types of evasion used by crickets – running and dodging – both adhere to the cybernetic “gain and maintain the preferred inter-animal distance” rule that is used by rats, despite the large differences in their body morphology and their mechanics of locomotion. Not only does this show that cybernetic rules can be applied to two different organisms, but the fact that crickets and rats come from vastly different evolutionary lineages supports the idea that cybernetic rules may be widely, if not universally applicable (Powers, 1973). This possibility has wide-ranging implications, both for understanding the behaviour of organisms, and for the development of artificial systems (e.g., robotics).

Chapter 4

SIMULATING THE SYSTEM

A further test of whether or not one has deduced the rules by which animals behave is to re-create the behaviour of the system by implementing those rules in a simulation. Although simulation is not necessarily a test of the mechanisms that underlie the rules, since it is possible to re-create a given behaviour by employing any number of vastly different architectures – as is evidenced by the fact that rats and crickets use the same rule, presumably generated by different sets of machinery – simulating the behaviour allows one to determine whether or not the rule itself is correct. Several computational models of PCT have been constructed previously. McPhail and Tucker (1990) demonstrated that variable behaviour, used to minimize the disparity between the reference value and perceived value of certain variables could re-create crowd behaviour better than methods based on specific behaviours in response to reward structures. Powers (2008) has created a variety of computational models, both hierarchical and non-hierarchical, that reproduce a range of behaviour. For example, at the organism level, he has reproduced the movement of *E. coli*, the following behaviour of geese, and crowd behaviour. At the sub-organism level, his hierarchical model of a human arm, using 14 control systems, demonstrates the ability not only to learn to co-ordinate the control systems on its own, but also to learn to perform complex patterns of movement.

In order to simulate the dodging behaviour, an agent-based modeling approach at the organism level was used that implemented the distance-regulation rule in freely-behaving agents within a simulation. In addition to the cybernetic rule, the behaviour of the agents was governed by some other basic movement rules (see below). The model that is discussed here is not a hierarchical model, nor does it have the capacity

for re-organization (i.e., learning). Despite its simplicity, and even though the agents were not explicitly coded to dodge, both the quantitative and qualitative aspects of the behaviour were re-created.

Method

Properties of the Simulation

The MASON library (Luke, Cioffi-Revilla, Panait, Sullivan, & Balan, 2005) in the Java programming language was used to create an agent-based simulation incorporating the “maintain a minimum distance” cybernetic rule that had previously been described in both rats and crickets during food defense behaviour. The simulation was comprised of three main types of classes – the *rat* class, the *arena* class, and the *food item* class, which governed the behaviour of each type of object. In addition, the *observer* class handled data collection, but did not participate in the simulation. Finally, the *agents simulation* class controlled the stepping of the agents within the simulation and the *graphical UI* class graphically displayed the simulation.

Because body shape is known to interact with environmental structure, influencing overt behaviour (e.g, May, Schank, & Joshi, 2006; May & Schank, 2009; May, Schank, & Joshi, 2011), the “rat” agents were drawn in a quasi-realistic style, with pointed noses, and larger posteriors, and the boundaries of the virtual arena were circular (see Figure 4.1). Although the algorithm used by the agents to determine their distance from one-another employed trigonometry based on the location of the centre of each agent, the food item, once one agent was in possession of it, was placed at the “nose” end of the defender so that, as was the case in the rats and crickets, the nose would be the region that would be both targeted by the robber and defended by the defender. And although one of the agents was assigned the properties of the defender, and the

other, the properties of the robber at the beginning of the simulation, the identity of the robber and defender, after their initial placement, switched, depending on which agent possessed the food item.

The simulation began with the random placement of a defender, a robber, and a food item inside the virtual arena. Initially, the robber would remain stationary, and the food would be tracked and picked up by the defender. Once this occurred, the robber would also begin tracking the food. After a certain number of timesteps, the food item would disappear, and a new food item would be randomly placed in the arena, which would then be tracked by both agents.

The movement of the rat agents at baseline was a correlated random walk, assuring that they could not turn in unrealistic directions (i.e., the movement of the agent at each time step was correlated with its position during the previous timestep). A degree of randomness (probability of a random movement at each timestep was 0.25) was included in the movement of the agents, by implementing the Mersenne twister pseudo-random number generator (Matsumoto & Nishimura, 1998), to ensure that the agents would not become stuck against walls or each other. Both the defender and robber were given the ability to locate the food with a certain probability, which got proportionally more precise as the distance to the food item was minimized.

Once one of the rat agents was in possession of the food item, which occurred when the agent was located at a particular distance to the centre of the food, that agent would take on the properties of the defender. The defender would remain in one position, moving only if the boundary of the minimum distance to itself and the robber were violated, at which point, it would again implement a correlated random walk until the boundary condition were satisfied. The robber, on the other hand, would track the food item with the same algorithm employed earlier, but with the addition of a motivation variable, which made tracking more accurate as it increased

– and the value of the motivation variable increased in a linear fashion, based on the number of time steps since the last robbing attempt of the robber. Once the robber was close enough to the food, it would make 5 attempts to steal the food, after which, its motivation variable would be reset to 0, and it would move away from the defender. Due, in part, to the imprecision of the collision-detection algorithm used by the rat agents, it was possible for the robber to steal the food item from the defender, although this rarely occurred.

Analysis

The simulation was run for 2000 timesteps. During this time, once the boundary condition of the dodging rat was violated, the observer class would create a new, sequentially numbered .csv file that recorded the location of the noses of both the robber and defender, the distance between the noses of the agents, and how far the agents moved during each timestep. The output file was ended when the robber's motivation variable reached 0. The graphical output of the simulation was simultaneously recorded to ensure that the behaviour being analyzed was consistent with the types of behaviours that were analyzed for the rats and crickets (i.e., that the dodges occurred away from the walls of the arena and that the robber was unsuccessful in stealing the food item) and to ensure that the collision detection algorithm was working properly so that the agents were not overlapping one-another. Of the 247 output files generated, 20 were deemed suitable for analysis.

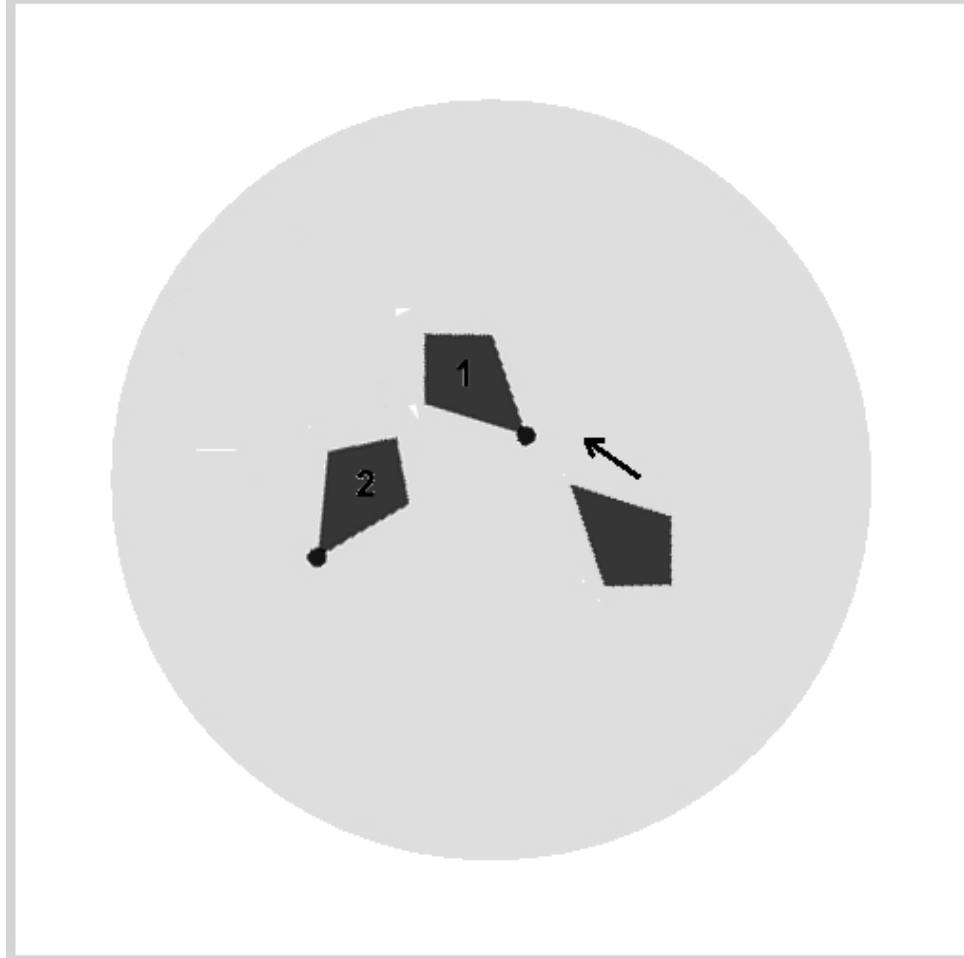


Figure 4.1: Simulated rats behaving in a circular arena. 1 is the position of the defender before the robber approaches, and 2 is the position of the defender after the approach of the robber.

Results

Correlations

As with the rats and crickets, the relationship between the movement of the robber and defender, measured as path length (in pixels) was examined to ensure that the system was functioning the way it was supposed to be (i.e., that the virtual robber and defender were compensating for each other's movements). And as with the animals, the movements of the virtual animals were correlated with one-another, $r_{18} = .808, p < .001$ (see Figure 4.2A).

The test for the controlled variable was also performed. The relationship between the movement of the robber and the inter-agent distance was analyzed. Because the output was stopped when the robber lost motivation (i.e., when the robber stopped robbing), this distance was used as the inter-agent distance, instead of, as had been done for the animals, the inter-agent distance between the position of the robber when it had stopped robbing, superimposed with the position of the defender when it stopped dodging. Regardless of this difference, there was no significant correlation between inter-agent-distance and robber path, $r_{18} = .288, p = 0.218$ (see Figure 4.2B). The higher magnitude of the correlation compared to those seen in the animals, nevertheless, probably owes both to the difference in the measurement used, as well as to the imprecision of the collision-detection between the agents.

Variability

Finally, in order to illustrate that the simulation behaved in the same way as the rats and crickets, the variability in movement of the robber and defender was contrasted with the variability in inter-agent distance, as measured by the coefficient of varia-

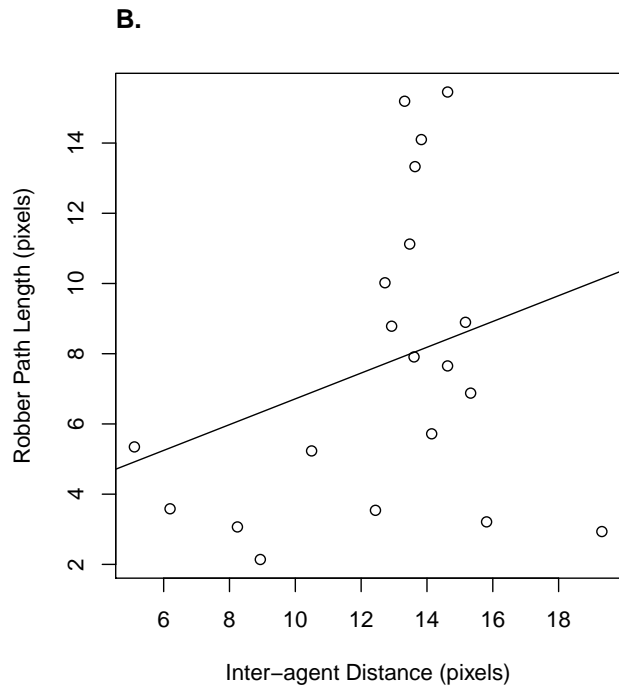
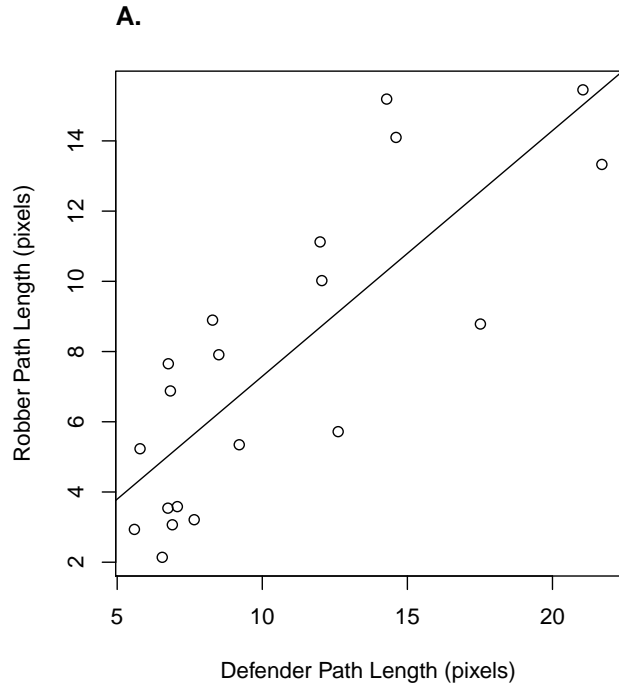


Figure 4.2: A. Correlation between robber path length and defender path length. B. Correlation between robber path length and inter-animal distance.

tion. As was seen in the animals, inter-agent distance was the least variable of the three parameters measured, and the specific value of the coefficient of variation was comparable to the values calculated using the data from the animals (see Table 4.1).

Discussion

The same pattern of results seen in both the crickets and rats holds for the simulation. Specifically, the correlation between the robber and defender movement indicates that the simulated agents also compensate for each other's movements. The test for the controlled variable, in which the behaviour of the presumed controlled variable (distance) is assessed with respect to perturbations of the system (the movement of the robber), demonstrated that these two measures were not correlated. The lack of a correlation confirms that the value of the distance parameter is being actively protected against perturbations by the defender. Further, the amount of variability seen in the distance parameter, as measured by coefficient of variation, was smaller than the variation in both robber and defender path length. This finding, as was also seen in both the rats and crickets, is evidence that inter-agent distance is what is "important" to the defender.

In some ways, these findings are trivial, given that the simulation was explicitly programmed to behave this way. However, the fact that the simulation behaves in the same way as the rats and crickets suggests strongly that the animals are, in fact, using this specific type of cybernetic distance-regulation rule in order to defend food. Although the operation of the simulation demonstrates the validity of the cybernetic rules discovered to be governing food protective behaviour in both rats and crickets, it is only an intermediate step in testing hypotheses concerning the generation of the rule used by the animals. The agent-based model, especially considering the

Table 4.1: Coefficients of Variation, $\frac{\sigma}{\mu}$

Measure	Dodging Crickets ¹	Running Crickets ²	Rats ³	Agents ⁴
Dodge Path Length	0.452	0.465	0.463	0.471
Rob Path Length	0.759	0.872	0.867	0.560
Distance	0.341	0.193	0.392	0.267

¹ $n = 56$

² $n = 84$

³ $n = 246$

⁴ $n = 20$

artificial way in which the agents detect food and each other, should not be used as a model for the mechanisms governing the implementation of the cybernetic rule. There is, however, a different kind of model that can be used to begin to address the question of mechanism. Neuromorphic engineering is, broadly speaking, a technique with which quasi- realistic neural hardware can be embedded in ecologically relevant robotic bodies (e.g., Krichmar & Wagatsuma, 2011). If the hardware were realistic enough, it may be possible to test hypotheses, in conjunction with evidence from experimental neurophysiology, as well as a behavioural approach similar to the one outlined in this thesis, about not only the rules that animals use to behave, but also *how* those rules are generated (see Chapter 6 for further discussion).

Chapter 5

PERTURBING THE SYSTEM: EXPLORING THE PARAMETERS OF THE DODGING SYSTEM IN RATS

Once it was determined that the same cybernetic, rather than S-R, rule governs the food protective behaviour of both rats and crickets (H. C. Bell & Pellis, 2011; H. C. Bell et al., 2012), the next question was one of how robust the dodging system is. If it were perturbed enough, could the system be destabilized to the point that it would begin to behave like an S-R system (but see the Discussion at the end of this chapter for an explanation of why these two types of system are not necessarily distinct)? If not, how would the parameters of the cybernetic system be changed by different manipulations?

There are many ways to disrupt a homeostatic system, and the approach taken depends on the type of question being asked, which falls into two main categories. The first type of question is concerned with *whether or not* the system is cybernetic, which can be addressed by driving the defended parameter of a maintained variable into a range outside of the set point of the system – the idea being to see whether or not the system is *truly* defending a specific parameter value of the variable of interest. This is what was done with the rats and crickets in the previous chapters, and occurred during the natural behaviour of the animals when the robber violated the minimum inter-animal distance that was being maintained by the defender.

The second way to disturb a cybernetic system, which is the focus of the present chapter, is to change some aspect of the operation of the feedback loop in order to determine how the functioning of the system as a whole is altered. For instance, by altering the sensitivity of the thermometer inside the thermostat controller, or the perhaps the speed with which the furnace can be turned on, one can imagine that the

gross behaviour of the system would be changed. Specifically, with adjustments of this sort, one would expect that the level of the controlled parameter would become more or less variable – that is, the amplitude of the oscillations of the controlled parameter around the set point of the system would change. Another modification that could be made is to adjust the set point of the system. Depending on how extreme the new set point was, the implications to the behaviour of the system might be quite unpredictable, given that there is likely an optimal range of set temperatures under which the thermostat is designed to operate. For example, if one were to set the temperature to 50 degrees Celsius, it is quite possible that the output of the furnace would never be sufficient to bring the ambient temperature of the room up to this level alone (i.e., in the absence of a similar ambient external temperature), and the system would be in a perpetually “on” state, obtaining a constant temperature, but not the desired temperature.

Biological systems, however, do not tend to operate like this. At some point, the physical demands of continually attempting to reach an unobtainable goal state are superseded by excessive perturbation of lower-order goals. In the thermostat system, there is only one level of control. In biological systems, there are presumably many (see e.g., Powers, 1973; Wiener, 1948), with lower-order, immediate goals (e.g., attaining satiation) needing to be satisfied before higher-order, long-term goals (e.g., finding a mate).

The aim of the current set of experiments was to explore the robustness of the function of the rat dodging system to alterations of the sensory and processing machinery that the animals use to sense and maintain inter-animal distance. One might reasonably assume that, at some level, in the absence of external information, the system might “default” to some other type (e.g., S-R), whereby the incoming information becomes so sparse that the organism is no longer able to detect, and therefore,

to compensate for changes to the defended parameters. The lack of compensation occurs because, in the absence of external information, the “best guess” of the system may simply be the information available at the beginning of the action.

In order to investigate the effects of a disturbed robbing and dodging system, two approaches were taken. For the first, two sensory modalities important to rats, sight and tactile information via vibrissae, were removed, hereafter referred to as, the *V* (vibrissae removed) and *VE* (vibrissae removed and enucleated) conditions (Pellis et al., 1996b). For the second, rats were raised in isolation, which is known to affect both brain development, and behaviour in social contexts, in ways that would be suspected to influence the performance of robbing and dodging behaviour, hereafter referred to as the *I* condition.

For all of the experimental manipulations, there are three broad classes of hypotheses, from which three sets of predictions can be derived (see Table 5.1). Because some of the predictions are the same for more than one hypothesis, what was of interest was the overall pattern of results. In order to get a better understanding of the overall results, the results for all three experiments are presented together, following the methods section.

For all three experiments, the video footage that was analysed was collected during the course of two previous studies. Experiments one and two involved re-analysing video footage from Pellis et al. (1996b), and experiment three was carried out on footage from Pellis et al. (1999). Re-analysis was done in lieu of performing new experiments because it was felt that the existing data could be used to address the previously-stated questions, in addition to not requiring the unnecessary use of further animals.

Table 5.1: Competing Predictions with respect to Three Hypotheses

	H_1^1	H_2^2	H_3^3
P_1	Distance remains uncorrelated with robber movement	Distance remains uncorrelated with robber movement	Distance may or may not be correlated with robber movement
P_2	Value of the distance parameter is the same as intact animals	Value of the distance parameter may be altered	Value of distance parameter may be altered
P_3	Distance is the least variable parameter	Distance is the least variable parameter	Path length is the least variable parameter
P_4	Correlation between robber and dodger movement as high as intact animals	Correlation between robber and dodger movement may be lower than intact animals, but still significant	Correlation between robber and dodger movement not significant

¹ Hypothesis 1: System behaves cybernetically; animals are able to fully compensate for experimental manipulation

² Hypothesis 2: System behaves cybernetically; animals are not able to compensate fully for experimental manipulation

³ Hypothesis 3: System behaves in S-R fashion; animals are not able to compensate for experimental manipulation at all

Experiments 1 and 2: General Method

For each experiment described below, a subset of the original data was analysed.

32 male Long-Evans hooded rats, born and raised at the University of Lethbridge, were used for the original study. Eight of the subjects were enucleated shortly after birth (see Experiment 2 below). Following weaning at 21 days of age, all subjects were housed in same-sex pairs with littermates in standard laboratory wire mesh cages (17 x 25 X 20cm) on 12:12 h light:dark cycles. Eucleated subjects were always housed with sighted cagemates. Food and water were provided *ad libitum* until shortly before the subjects were tested in the dodging paradigm.

The pairs of subjects were housed together continually until 80 days of age. At the beginning of the food deprivation period, the rats weighed 200-250g, and were then subjected to a partial food deprivation schedule. Their initial weight was reduced to, and then maintained at 80-85% of their original mass.

Tests were conducted between 1200 and 1300 h, with pairs of rats placed in a thin Plexiglas cylinder, 40 cm in diameter and 45 cm high. It was placed on a table with a clear glass top. An inclined mirror, through which the rats could be viewed and videotaped, was located underneath the table. One rat was given a piece of food pellet (Purina Rodent Chow) of approximately 2.5 g. The behavior of the other rat attempting to take the food pellet was then videotaped. After at least 10 trials in which one rat attempted to steal the pellet and the other dodged away were videotaped, the other rat was given a food pellet. For this test, the vibrissae of both animals were closely cropped. (see Pellis et al., 1996b).

Experiment 1: Disrupting Sensory Input I

Method

For the first experiment, the dodging behaviour of animals that had had their vibrissae shaved off were compared against completely intact animals. Given that this manipulation had the least effect on dodging behaviour in the original study, it was hypothesized that either the dodging system would remain completely intact, or that there would be some slight modifications to the parameters of the system, but that overall, the cybernetic nature of the behaviour would be preserved.

Analysis

Footage from six pairs of animals was used. Six dodges for each member of each pair was analysed, except for one animal in pair number 2, for which not enough suitable dodges could be analysed. This animal's data were not included in the analysis. The dodging behaviour of these animals was compared against dodges from five pairs of completely intact males, which was collected in the course of the original research that is analysed here in Experiment 3, which was done because labeling issues prevented definitive identification of the relevant control group used for this study. It should be noted that the same set of control animals was also used for Experiment 3.

As with the rats and crickets before, only dodges that occurred away from the edges of the arena were selected, and each dodge was analysed using the digital tracking system, Vicon Motus (see Chapters 2 and 3).

Experiment 2: Disrupting Sensory Input II

During the course of the original study, the combination of enucleation with the removal of the vibrissae produced the largest change in dodging behaviour (Pellis et al., 1996b). It was therefore hypothesized that this manipulation would be most likely to destabilize the cybernetic dodging rule to the point that the rats would begin to behave in S-R fashion, and so this was the experimental group that was analysed next.

Method

Half of the subjects were enucleated (see below) and had their vibrissae removed via shaving. The other half of the subjects, which were littermates of the enucleated half, were left completely intact.

Surgery

The subjects were enucleated bilaterally under cryoanesthesia within 8h of birth. Upon removal of the eyes, the sockets were packed with Gelform (Upjohn Co.), and one suture was used per eye to keep the wounds closed. The intact littermates that were used as controls for this manipulation were also anesthetized, but left otherwise intact (see Pellis et al., 1996b).

Analysis

12 male Long-Evans hooded rats (6 pairs) were used. 4 video clips were taken for each subject, as well as for its normal partner. For this experiment, the normal cagemate (and dodging partner) of the enucleated rats were used as the control group,

because labeling issues with the footage prevented precise identification of the control versus experimental animals, which could not be resolved by behavioural observation alone. Nonetheless, for all of the pairs, one of the animals was marked, and the other was not, which was the original method of identifying their respective experimental conditions. Therefore, following the re-analysis of the data, group identification was made by comparing the values of the distance between the animals at the start of the dodge against those published in the original research.

Experiment 3: Disrupting Information Processing

It is known that raising rats in isolation changes several aspects of their behaviour (Jones, Marsden, & Robbins, 2001), including some aspects of dodging behaviour (Pellis et al., 1999). What seems to be the crucial experience lacking when rats are isolated during the juvenile period is that they are unable to engage in rough and tumble play with other juveniles. Rats raised with the ability to interact with – but not to play with conspecifics produces many of the same deficits in social behaviour as rats reared without any social contact at all (Pellis & Pellis, 2006). In addition, when the brains of rats raised in conditions without access to other juveniles are analyzed, alterations are seen in the neuronal morphology in both the medial prefrontal and orbitofrontal cortices (H. C. Bell, Pellis, & Kolb, 2010).

Method

Subjects

Eleven pairs of male Long-Evans rats, born and raised at the University of Lethbridge, were used. Six pairs of subjects, those in the *I* condition, were reared from weaning

at 21 days of age in divided wire mesh cages (total dimensions: 17 x 25 x 20 cm). Raising rats in divided cages removes the experience of juvenile peer-peer play, which is crucial to brain as well as adult social behavioural development (H. C. Bell et al., 2010). Subjects were maintained on a 12:12h light:dark schedule, and were given food and water *ad libitum*. At 90 days of age, subjects were reverse-housed for 35 days, and then were housed in non-divided cages for an additional 15 days, which has been shown to ameliorate the acute, but not long-term effects of isolation-housing (Potegal & Eison, 1989). 2-10 days prior to testing, subjects were food-restricted, such that their weight was reduced to about 80-85% of their free-feeding weight. The rearing conditions for the controls was identical, except that the cages were not divided.

Testing Procedure

Trials were conducted in a manner identical to, and using the same experimental setup as those carried out for Experiment 1 and 2, with only a few minor differences. The time of testing was between 1400 and 1600h, and the food was removed after five, not 10 dodges per animal. The animals were tested against their similarly isolation-raised cagemates.

Analysis

A minimum of seven dodges were analysed for each pair. As with all previous work, only dodges that occurred away from the centre of the arena were selected. The dodges of the isolation-raised animals were compared against the dodges of the same intact controls that were used for Experiment 1. These animals were part of the original research conducted for the present analysis.

Results

Correlations

The correlation between robber movement (i.e., path length) and inter-animal distance was tested. For all conditions, the correlation between these two measures was not significant ($p > 0.05$), with the strongest correlation, $r = -0.097$, occurring in the *VE* condition (see Figure 5.1). This lack of a correlation indicates that in all instances, inter-animal distance was being controlled by the subjects.

In order to determine whether or not the animals were still compensating to the same degree for each other's movements, the relationship between the robber and defender movement was tested. In all conditions, this correlation was significant. The highest correlation, $r = 0.824$, was seen in the *V* animals, whereas the correlation for the controls (for the control group used in Experiments 1 and 3) was $r = 0.802$. For the *VE* and *I* animals, the correlations were $r = 0.767$ and $r = 0.742$ respectively (see Figure 5.2). These correlations indicate that in all conditions, the animals continued to compensate for one another's movements, although the compensation was slightly reduced in the *VE* and *I* animals.

Variability

Once again, the variability of inter-animal distance, as well as of the path length of both the robber and defender were analysed using coefficients of variation. Under all conditions, the same pattern of results was observed: robber and dodger path lengths were roughly equivalent in terms of variability; whereas inter-animal distance was far less variable than either of the path length measures. This demonstrates that what is of importance to the behaving rats is not the particular distance being moved when

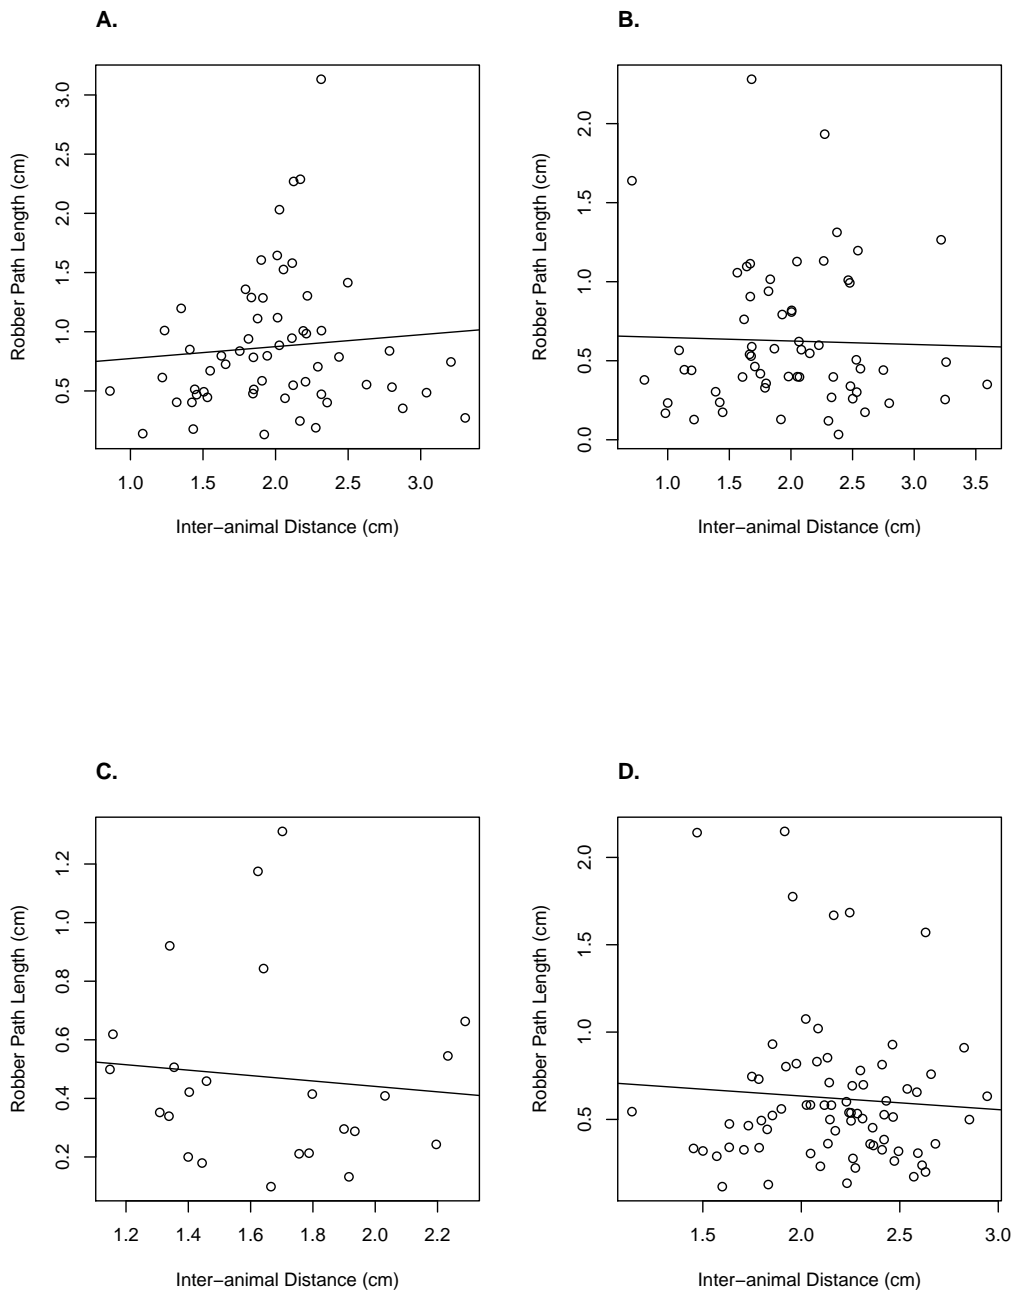


Figure 5.1: Robber path length vs. inter-animal distance for A. Controls, B. Vibrissae Removed, C. Vibrissae Removed and Enucleated, and D. Isolates

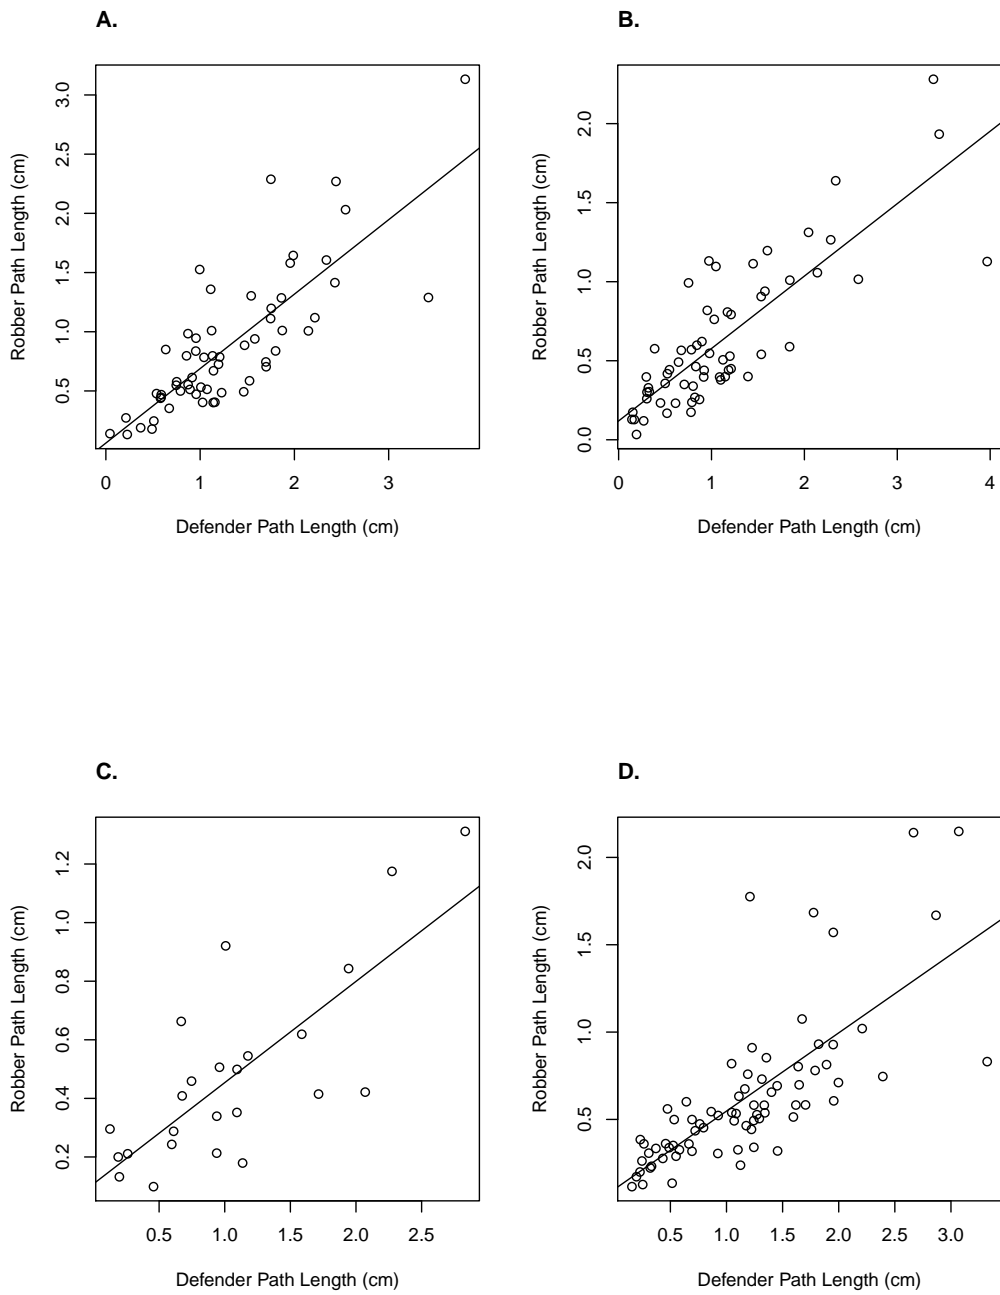


Figure 5.2: Robber path length vs. Defender path length for A. Controls, B. Vibrissae Removed, C. Vibrissae Removed and Enucleated, and D. Isolates

protecting against robbing attempts, but the distance from the snout of the dodger to the snout of the robber (see Table 5.2).

Orientation

The final orientation of the defender was examined. In all conditions, chi-square analyses revealed that the experimental groups did not differ significantly from its control group in terms of final defender orientation ($p > 0.05$). In addition, as was seen in the crickets, facing directly away from the robber was the preferred final orientation for all groups (see Table 5.3).

Differences Between Conditions

Between-subjects t -tests were run on the parameters measured for each condition against its relevant control group (see Table 5.4).

For the V animals, Welch's two-sample t -tests were run using fully intact animals as the comparison group. In the experimental group, both the robber and the dodger in each pair had their whiskers shaved. The results of the analysis revealed that neither the distance at the initiation of the dodge, nor the minimum maintained distance between the animals differed between the groups ($p > 0.05$); however, statistically significant decreases were seen in the experimental animals for both robber and dodger maximum instantaneous velocity ($t_{119.369} = 2.480, p < 0.05$ and $t_{119.184} = 2.141, p < 0.05$ respectively), as well as a significant decrease in robber path length ($t_{111.126} = 2.645, p < 0.05$). The mean path length of the dodgers in the experimental condition was also lower than that of controls; however, this difference was not statistically significant ($p > 0.05$).

Welch's t -tests were also run for the EV group, with the intact cagemates that

Table 5.2: Coefficients of Variation, $\frac{\sigma}{\mu}$

Measure	Controls ¹	V ²	VE ³	I ⁴
Dodge Path Length	0.571	0.728	0.667	0.604
Rob Path Length	0.663	0.717	0.671	0.684
Distance	0.253	0.302	0.201	0.168

¹ $n = 60$, intact controls used in Experiment 1 and 3

² $n = 62$

³ $n = 24$

⁴ $n = 75$

Table 5.3: Frequencies of Final Defender Orientation by Experimental Group and Corresponding χ^2 Values

	Facing Away	Other Orientation	χ_1^2	p -value
<i>V</i>	41	21	6.461	< 0.05
<i>VE</i>	18	6	6.001	< 0.05
<i>I</i>	62	13	32.012	< 0.05

they were dodging against used at the control group. Both the distance at the initiation of the dodge, as well as the minimum maintained distance were reduced in the experimental animals ($t_{35.159} = 2.760, p < 0.05$ and $t_{31.186} = 4.075, p < 0.05$ respectively). On all other measures, experimental animals did not differ from controls ($p > 0.05$).

The same analyses were conducted on the *I* animals, using the same controls as for the *V* animals. Although the distance between the animals at the initiation of the dodge did not differ between experimental and control animals ($p > 0.05$), the minimum maintained distance was greater in the experimental animals than in controls ($t_{103.861} = 2.122, p < 0.05$). In addition, both the maximum instantaneous velocity and the path length of the robber were lower in the experimental animals ($t_{118.947} = 2.504, p < 0.05$ and $t_{105.296} = 2.814, p < 0.05$ respectively). Mean maximum instantaneous velocity, as well as mean path length for the dodger were also lower for the experimental animals, but the difference was not significant ($p > 0.05$).

Discussion

It was predicted that, given the importance of eating to survival, that dodging should be fairly robust to minor disturbances, but that more major disruptions would alter at least some parameters of the system, or even drive the system to become S-R – although the distinction is somewhat trivial, because a cybernetic system, in the absence of enough sensory information for the organism to compensate for the actions of conspecifics, and therefore appearing to be S-R, could still be fundamentally cybernetic. There need not be any kind of switch to a different kind of rule or way of processing incoming information about the environment. The cybernetic rule would continue to exist, but the organism would be incapable of instantiating it, because it

Table 5.4: Overall Changes to Dodging Parameters, Relative to Controls

	V	VE	I
Initial Distance	–	↓	–
Maintained Distance	–	↓	↑
Dodger Path Length	–	–	–
Robber Path Length	↓	–	↓
Dodger Max Velocity	↓	–	–
Robber Max Velocity	↓	–	↓

would be incapable of compensation. In fact, Rosenblueth et al. (1943) speculated that some behaviours, such as a frog striking at a fly, are essentially S-R, because they occur so quickly that adjustments made based on feedback once the movements have been initiated are not possible. But even a striking behaviour can, in some larger sense, be viewed as cybernetic – because the animal will surely try again if it misses (and the prey is still there).

This series of experiments demonstrates that the dodging system in rats is relatively robust to manipulations. For all manipulations examined, the general pattern of results supports Hypothesis 2: that the dodging system is able to compensate for the altered functioning of some elements of the feedback loop, but that some of the parameters of the system are changed (see Table 5.1). Crucially, all of the cybernetic components of the system remained intact for all of the groups of animals.

The different groups did display differing kinds of parameter shifts, however. In the *V* condition, although the inter-animal distance did not change, the maximum velocity of both robber and dodger was decreased. Likely not unrelated to this, the path length of the robber was also shorter than that of controls, with dodger path length demonstrating the same tendency, although the results were not statistically significant.

In the *VE* animals, alterations were seen in both the distance at which dodges were initiated, and the maintained minimum distance, which were both smaller than in controls. It is suspected that this difference is due to the relative inability of the *VE* rats to detect conspecifics.

Finally, with respect to the *I* rats, there was an increase in the maintained inter-animal distance. It seems as though the isolation-reared rats became hyper-sensitive to the approach of conspecifics, possibly because they lacked experience interacting with others that would allow for a more nuanced reaction. Interestingly, a similar

effect on inter-animal distance arises from domestication. In the play fighting of rats, the distance at which a defender begins to react to an attacker during play fighting differs, with wild rats reacting at twice the distance of domesticated rats (Himmler et al., in press). Additionally, a decrease in both robber velocity and path length were seen (keeping in mind that both the robbers and dodgers in the *I* condition were isolates), possibly reflecting some diminished motivation to pursue the dodger.

Chapter 6

CONCLUSION

Is Dodging in Rats Cybernetic?

The vast majority of behavioural descriptions in the literature are formulated in strictly S-R terms. Because of this, the focus of the present research has been to test, explicitly, whether or not some behaviours, previously thought of in terms of S-R, can be re-defined in terms of, and better understood as arising from, cybernetic mechanisms. At the outset, it was decided to investigate a behaviour that had previously been described as an S-R system, robbing and dodging in rats. Until this work, it had been thought that the magnitude of the two-dimensional angle swept out by the dodger (viewed from above or below) during the pivot was directly related to contextual information, such as the quality of food being consumed, as well as relative partner identity (i.e., dominant or subordinate) (Whishaw & Gorny, 1994; Pellis et al., 2006).

However, given that rats are able to learn, if relatively fixed dodging angles were being used by the dodger once the dodge were initiated, the dodger would not change its behaviour, regardless of the subsequent behaviour of the robber the robber would eventually learn this, and so it would seem reasonable that the robber would be able to defeat the dodger. In fact, robbers rarely defeat dodgers, implying that, either robbers are unable to learn, or that the description of the system was incorrect.

A simple way for the dodger to solve, continuously, the problem of the robber's food-stealing attempts would be to maintain a minimum distance between itself (specifically, the front of itself, where the food is located) and the stealing end (front) of the robber. This is a cybernetic rule, whereby the dodger, in real time, controls

the distance between itself and the robber, compensating for the movements of the robber. It should be noted that the robber might also be using a cybernetic rule to rob, but instead of maintaining its distance to the food, its goal would be to reduce its distance to the food to zero.

In order to test the cybernetic hypothesis against the original S-R description of robbing and dodging, the variability in the magnitude of the dodging angle was compared to the variability in the inter-animal distance. If it were the case that dodgers were behaving in a S-R fashion, it would be expected that the dodging angle would remain relatively constant when contextual information remains the same, whereas the inter-animal distance should be more variable. If, on the other hand, what was important to the dodger was inter-animal distance, it would be expected that it would remain relatively constant, and that the dodging angle should be more variable. After digitally tracking the dodging behaviour of seven focal animals, each paired with both a dominant and a subordinate partner, and each given two different types of food, it was found that, in any given context, the dodging angle (and, in fact, every other parameter that was measured) was much more variable than the inter-animal distance. This result was true both across and within individual animals (see Chapter 2 and H. C. Bell & Pellis, 2011).

The second part of the analysis was to determine whether inter-animal distance were truly a controlled parameter. If a system is cybernetic, and is actively working to control a particular parameter, then, as the system is perturbed – as the value of the controlled parameter is altered – the system should compensate in order to return the parameter to its preferred level. A consequence of this behaviour is that the level of the parameter should be uncorrelated with disturbances to the system. In the case of the rats, the robber perturbs the value of the controlled parameter (distance) by attempting to get close enough to the dodger to steal its food, reducing the inter-

animal distance to a level below the dodger's preferred minimum. Thus, if distance is being controlled by the dodger, one would expect that inter-animal distance should not be correlated with the movement of the robber – in this case, the angle swept out by the robber during its robbing attempt. The relationship between inter-animal distance and robber angle was compared, and no correlation was found (see Chapter 2). However, in order to ensure that the system was operating the way it was thought to be, the movement of the robber was compared to that of the dodger's, which should be correlated if the movements were related to one-another, which they should be if the animals were compensating for each other's movements. There was, indeed, a significant correlation between robber and dodger angle (see Chapter 2).

Finally, it was found that, in fact, contextual information did influence dodging decisions – but not in the way previously described. The inter-animal distance at the onset of the dodge varied with respect to two types of contextual information, food type and partner identity. The preferred minimum distance varied as a function of partner identity, but not food type, suggesting that some contextual information influences the parameters of the underlying cybernetic behaviour of the dodger (i.e., in some situations, it might be prudent for the dodger to maintain a greater minimum distance, whereas in others, a smaller distance will suffice).

How widespread is the cybernetic dodging rule?

If it is the case that rats use a cybernetic rule when protecting food, and given that food protection is a ubiquitous problem for organisms, one might ask if other species implement similar types of rules when engaging in similar behaviour – and just how far away, in phylogenetic terms, might these types of rules continue to be found? Fortunately, another group of animals, separated from rats by about 500 million

years of evolutionary history, also use dodging to protect food: crickets.

Some have argued that crickets, and arthropods in general, have nervous systems so simple that they must be S-R machines, and that their behaviour is essentially ballistic (Schöne, H., 1984; W. J. Bell, 1991). In other words, that the nervous system of crickets is so simple that they are unable to process information quickly enough to be able to update their behaviour once it is initiated; thus they rely only on relatively stable sets of information, such as contextual cues, to behave. Again, the two competing hypotheses were pitted against one-another.

The analysis began by measuring variability in inter-animal distance versus, in the case of the crickets, the path length of the dodger. Path length was measured because, unlike rats, which use pivoting most of the time (and only the instances in which pivoting was performed were evaluated) to evade a robber, crickets are just as likely to run straight away from a robber as they are to pivot, depending on whether the robber approaches the dodger from the front or the rear. If the front is approached, the dodger pivots; however, the dodger will run straight away if approached from the rear. The crucial point is that, whether the dodge angle or path length is used, it is a proxy measure for the movement of the dodger, and should be relatively constant, if behaviour is essentially ballistic. It should be noted that the rat data were re-analysed using path lengths instead of angles, and that the results were unchanged. As was the case with the rats, inter-animal distance was significantly less variable than any other parameter measured (see Table 3.2) and was not correlated with the movement of the robber, and the movement of the robber was correlated with the movement of the dodger (see Chapter 3 and H. C. Bell et al., 2012). It was concluded, therefore, that crickets use the same cybernetic distance regulation rule that rats use to protect food from conspecifics.

Wiener's approach: Can we simulate the behaviour?

One way to test whether or not one has correctly understood the behaviour of an animal is to simulate that behaviour using the experimentally-derived rules, to see if the simulation behaves in the same way as the real organism (Schank & Alberts, 1997). I used MASON, a library written in Java (Luke et al., 2005), to create an agent-based model of the rat dodging system (see Chapter 4). The primary advantage of using an agent-based approach is that it allows the different agents in the simulation to be heterogeneous - that is, to possess different properties - as opposed to representative agent models, in which all agents of the same type are given identical properties. This is a key point, because the goal of the robber is to steal the food from the dodger, whereas the goal of the dodger is to evade the robbing attempt - that is, the goals of the two agents are not the same.

The dodger was given one rule: once it possessed a piece of food, it was to maintain a minimum distance between itself and the robber (i.e., if the robber were too close, do something until the robber is no longer too close). With respect to the robber, its only rule was to try to get as close to the food as possible. However, the robber's rule was modulated, as the robber was also instilled with a varying degree of motivation to track the food, whereby, after a few attempts to steal the food, it would move away from the dodger. With these two cybernetic rules, the dodging behaviour seen in actual rats and crickets was reliably replicated. In this simulation, it was even seen that if the robber approached the front of the dodger, the dodger would pivot away, but that the dodger would move off in a straight line if approached from behind, which was the same behaviour observed in the crickets (see Chapter 4).

How robust is the cybernetic system?

In order to elucidate the properties of the cybernetic food defensive system further, the robustness of the system to disturbance was analysed. This was done in two main ways: by disrupting the sensory inputs to the dodger, and by altering the way in which sensory input was processed in the brain of the dodger in rats. To begin with, the dodger's ability to use its tactile sense was impaired by shaving off its whiskers. This was done because it is known that rats use their whiskers to perform a variety of tasks (e.g., Bermejo, Harvey, Gao, & Zeigler, 1996; Carvell & Simons, 1995), and that removing the vibrissae modifies performance on those tasks. The results for the dodgers with shaved whiskers were statistically different from intact controls in only minor ways (see Chapter 5). Because the rats were behaving under lighted conditions, it is likely that their vision was sufficient to produce the behaviour, even in the absence of some tactile input. The second experiment involved disrupting both the visual and tactile inputs of the dodger. The analysis demonstrated that one of the main parameters of the dodging behaviour, the maintained inter-animal distance, was diminished. In addition, the distance at which the dodge was initiated was also diminished.

In the third experiment, control over sensory input was altered indirectly by changing the development of the brain during the juvenile period. When the behaviour of dodgers that were raised in isolation was analysed, it was found that all major aspects of the dodges remained the same as in the controls, except that the preferred minimum distance was increased (see Chapter 5).

Regardless of the manipulation, the dodging system of all of the rats continued to adhere to the cybernetic distance-regulation rule (see Chapter 5), demonstrating the relative robustness of this system to interference.

What if the Controlled Variable is Unknown?

Unlike the case of the rats and crickets, there are many instances in which it is not necessarily obvious what, if any, variable is being controlled by the organism. Also, unlike dodging, there are many behaviours that are not obviously tracking behaviours. Could these other types of behaviour be under cybernetic control? And if so, how can we narrow down the possible controlled variables in the system?

One possible approach involves tracking the movements of animals with software, analyzing every possible parameter, and then determining which parameters from that list remain the most stable. Although this certainly could be an effective strategy, one could also imagine that this would be a rather arduous task, generating a lot of data, and it is certainly possible that the controlled parameter might be missed. An alternative approach is to use a technique that focuses the researcher's attention on features of the animal's behaviour that are maintained constant relative to its partner. Such a technique is the Eshkol-Wachman Movement Notation (EWMN) (e.g., Golani, 1976; Pellis & Bell, 2011).

EWMN enables the experimenter to track the movements of parts of organisms, and of organisms as a whole, in relation to their other parts, to other organisms, and to aspects of the environment (Eshkol & Wachman, 1958). The resulting score, not unlike musical notation, is comprised of specialized symbols, which allows a trained reader to re-create the animals movements without having seen the behaving animals. By juxtaposing the animals movements in different frames of reference, with respect to their own body movements, to those of their partner and to the environment, one can identify constancies in the behaviour of the animals (e.g., Golani, 1976; Moran et al., 1981; Pellis, 1982).

For example, in the combat of male sage-grouse (*Centrocercus urophasianus*) the

birds approach one another face-to-face, then stand shoulder-to-shoulder and maintain this orientation, and from this position, they attempt to hit each other, on the head, with wing strikes (Wiley, 1973). An EWMN analysis showed that, this shoulder-to-shoulder also afforded the birds the opportunity to defend against these wing strikes, and so the birds maintained the optimal orientation to avoid being hit, while simultaneously, positioning themselves to launch an attack. This combined attack and defense maneuvering involved one bird performing compensatory movements to block the other bird from gaining the most advantageous position from which to strike (Pellis et al., 2013). The EWMN analysis identified that the birds actively maintained this inter-animal orientation, despite both birds moving backwards, forwards, sideways and in circles. Once candidates for constancies have been identified, other methodologies can be employed to test whether or not they actually exist.

In order to test this approach, the combat behaviour of male Madagascan hissing cockroaches (*Gramphadorhina portentosa*) was examined in a small pilot study. Other researchers have previously described combat behaviour in cockroaches as consisting of a suite of distinct behaviours (e.g., Clark & Moore, 1994; Clark, Beshear, & Moore, 1995). Given that during combat, the cockroaches often flip each other over onto their backs, it was postulated that the distinct behaviours observed may not actually be distinct, but rather, arise as a by-product of both animals simultaneously trying to flip one another over. The results of the study suggested that some behaviours observed during combat are by-products of a cybernetic rule that is being used simultaneously by the animals (see Appendix).

Where to go From Here? The Role of Neuromorphic Engineering

Neuromorphic engineering, pioneered by Carver Mead, originally focused on the implementation of large scale analog, as opposed to digital, circuitry in the development of robotics and AI (Mead, 1990). The notion was that, because biological systems are far better at information processing than artificial systems, perhaps mimicking their function, which is analog, would lead to significant advances in AI. Since its inception, neuromorphic engineering has gradually come to be associated with broader aims of biomimicry, incorporating ideas from embodied cognition such that realistic neural hardware is embedded in biologically plausible morphologies, which then act in the real world (e.g., Bernardet, Bermudez i Badia, & Verschure, 2012; Yang, Cameron, Lewinger, Webb, & Alan, 2012).

Ultimately, a complete research program would involve the demonstration that the control system properties of organisms can actually be instantiated in artificial models with as close to the same properties as biological systems as possible, and there are several reasons why simply simulating behaviour using methods such as agent based modeling are not fully satisfactory. These include the fact that robots do not require physical forces in the environment to be simulated, important aspects of which can be missed when creating virtual environments (Tamburrini & Datteri, 2005). By being required to act in the real world, insights are gained into how physical forces and the structure of the environment interact with the robotic agent to produce behaviour. For example, even in as abstract a situation as the Prisoner's Dilemma, robotic agents were able to offer more insight into this behaviour than were simulated agents (Grimaldi, 2012).

Given the neuromorphic research program, arguments could certainly be made that this has already been done for a number of systems (e.g., Bernardet et al., 2012;

Webb, 2008), but again, the question of why these insights have largely been ignored in the biological arena arises. Although it is always a risk that, even though the behaviour of an organism can be reliably re-created, the details of the mechanisms of the model are incorrect (Webb, 2000). However, if a cybernetic rule can be enacted in an entity with both realistic body morphology, as well as neuronal hardware, it does at least suggest that the particular mechanism is plausible as an explanation for the behaviour.

The Role of Embodied Cognition in the Production of Behaviour

Although the behaviour of organisms can be highly variable in the service of controlling a relevant variable, it is also important to realize that there are limits to how much variability is possible. Understanding the behaviour of organisms also requires that one be aware, in addition to how neural processes produce behaviour, the physical constraints on behaviour - that is, that the shape of the organism, its biomechanics, how its nervous system is structured, and the physical characteristics of its environment, all ultimately contribute to observable behaviour (Barrett, 2011).

Early authors, such as Wentworth Thompson (1917), recognized the importance of physics as a driving force, as much as natural selection, in the development of the morphology of organisms. But it seems a much more recent development that the role of environmental and morphological constraints in shaping behaviour have been investigated (e.g., Barrett, 2011; Pfeifer & Bongard, 2007). For example, the morphology of simulated agents has been demonstrated to affect their behaviour so much that, for instance, differing numbers of body segments produce vastly different behaviour, even though the algorithms governing behaviour are the same (May et al., 2011). The degree of bilateral symmetry in artificial agents has also been

shown to influence their locomotive efficiency in a virtual environment (Bongard & Paul, 2000). Even morphology of whiskers in both robots and simulated agents was demonstrated to alter the ability of robots and agents to follow walls in biologically realistic ways (Fend, Bovet, & Pfeifer, 2006). Ant cemetery building can be replicated using computer simulations by appealing to one decision rule, coupled with environmental structure (Martin, Bastien, & Albuquerque, 2002). Briefly, if an ant sees a corpse, it picks it up. The ant then moves off in a random direction. When the ant carrying the corpse encounters another corpse, the ant drops the corpse it is carrying. Using this model, cemetery formation tends to occur against walls – as is also the case with “real” ants. This is because the movement of the ants is restricted in these locations, and clusters, once they are formed, are more resistant to being dissipated. It is important to note that cemetery formation will occur whether there is just one ant, or thousands. The huddling of neonatal rat pups can similarly be explained by the interaction between “hardwired” thigmotaxis, body morphology, and environmental structure (Schank & Alberts, 1997; Schank, May, & Joshi, 2004; May & Schank, 2009). Two basic parameters explain huddling behaviour: movement (expressed as a probability dependent on previous movement) and thigmotaxis – the preference by the rat pup to be in contact with another object. As with the ant cemetery-building behaviour, aggregations of rat pups tend to occur where there are environmental irregularities (i.e., walls, corners) that restrict movement.

How is Variable Behaviour Generated?

One of the fundamental aspects of living organisms is that they are able to vary their behaviour in order to achieve goals, but at some level, it is not very helpful to state simply that variable behaviour is produced in the service of achieving some goal

state. Then the question becomes, regardless of what the goal of the animal is, how is variable behaviour created to begin with?

Ashby (1958) described his Law of Requisite Variety such that the larger the number of possible behavioural states of a control system, the larger the number of disturbances for which the system can compensate. In the case of the thermostat, there is only one possible disturbance to the system, which when that the temperature is too high, and therefore, the control system does not need to exhibit much variability in order to compensate. In contrast to the thermostat system, most living organisms are subject to disturbances to relevant variables that are orders of magnitude greater in number, in addition to the number of variables themselves being much greater. It is therefore necessary that a wide variety of control system behaviour be available in order for the organism to compensate. And it is known that organisms are capable of implementing highly variable strategies in order to control relevant parameters (Golani, 1976; Pellis & Bell, 2011).

In addition, most organisms have multi-modal sensory equipment, which can detect graded differences in stimuli (analog), unlike the thermostat, which is binary – both in its sensory capacity, as well as its behavioural options. This variability may simply be the result of the exact state of the system at that moment, which is different from the state of the system at previous time points, even under the same external conditions (i.e., slightly different initial network states, including the embedded nature of the neural hardware in the rest of the system, any parameter of which, when varying slightly, can lead to vastly different final states), or it may be the case that behavioural variability is a fundamental component of the system – and perhaps the sensitivity to initial conditions of the system produces variability that is, for all practical purposes, irreducible. Further, it is possible that organisms somehow co-opt and amplify environmental variability (e.g., quantum effects) in order to produce vari-

able behaviour. These possible mechanisms are not necessarily mutually exclusive. Regardless, the ability of biological networks to generate spontaneous activity has been well documented (e.g., Mazzoni et al., 2007). Variability, therefore, need not be directly selected for, although it may be adaptive, and may simply arise because of the chaotic properties of the behaviour of networks, which seems to occur in even very small systems (e.g., Sabarathinam et al., 2013).

However, it seems likely that fundamentally variable behaviour should be important for organisms, in that it decreases, for example, the ability of other potential predator (or prey) organisms of learning to predict the behaviour of the focal organism accurately. For example, many prey animals have protean movements in their escape behaviour, making it difficult for predators to keep tracking them (Driver & Humphries, 1988). There is some evidence to suggest that variability is actively generated by nervous systems, rather than being a by-product (Beck, Ma, Pitkow, Latham, & Pouget, 2012; Brembs, 2011). The fact that variable behaviour can be reinforced in numerous species further supports the idea that variability is not simply the result of noisy inputs (Neuringer, 2004). Some argue that variability is necessary in order for the organism to learn about its environment effectively, called world learning. The reasoning is that variability, at least partially, allows the organism to “test” its environment (Brembs, 2011).

In fact, when the behaviour of organisms is simulated, it seems to be the case that an element of randomness (i.e., variable behaviour) needs to be included so that the agents do not, for instance, become trapped in corners (e.g., Powers, 2008). Even in *E. coli*, the movement of which is governed by the mechanical action of molecules, such as glucose, binding to receptors which activate its flagella, the bacterium will tumble randomly in the absence of bound molecules, which proves to be an effective strategy to locate a food source eventually (Koshland, 1980).

General Conclusion

It is possible to view a cybernetic system in S-R terms – that is, to cut the loop so that only an immediate input and output are considered. This approach, however, fails to capture the dynamic nature of the system, and the fact that the organism's own behaviour alters the parameter(s) of the variable(s) being controlled, which affects its subsequent behaviour.

Cybernetic systems, if the behaviour of organisms can truly be described in those terms, are not passive recipients of stimuli, but active participants in how stimuli are experienced. It is the difference between how a rock and a human behave when pushed. The rock can do nothing to counteract the applied force – that is, the behaviour of a rock is truly a sum of the forces that act on it. In contrast, the human actively compensates to try to maintain an upright position. And in order to compensate, the behaviour of the human, in terms of what muscle groups are activated when, is highly variable, depending on factors such as the topographical structure of the environment (is the ground he is standing on at an angle?), the speed with which he reacts (is he fully alert and anticipating this, or half asleep?), learning effects (has this happened before?), and genetic/epigenetic effects (how quickly and effectively is his brain and body able to process and compensate?), and perhaps an element of randomness – either intentionally or unintentionally generated by the interaction of all those parts – at any given moment.

Organisms are not inert objects, at the whims of the forces around them. They behave like machines that can act to alter both their internal and external environments – although some argue that even the machine metaphor is not general enough to capture the qualities of life entirely (e.g., Rosen, 1991).

Variability in behaviour, even if the root causes of it remain unknown, does not

render the behaviour of organisms completely, or even largely, unintelligible but it may be the case that precise, moment-to-moment behavioural prediction is not possible. By appealing to a closed-loop conception of organisms, the understanding of behaviour can, instead, be approached on a different level. Instead of asking how a particular experimental manipulation alters the subsequent behaviour of an organism, one might instead ask how an experimental manipulation alters the parameters of the system. This is a subtly different question, but the difference is important, and requires that the parameters of the system be understood to begin with. Understanding what variables organisms may be controlling necessitates that organisms be understood on their own terms before they are used as model systems to answer larger questions.

The approach outlined in this thesis is a collection of methods by which these problems can begin to be addressed. Such an approach can be particularly useful for comparisons across species and contexts, as it provides a firm theoretical and methodological framework within which to identify similarities and differences (e.g., Finley, Ireton, Schleidt, & Thompson, 1983; Pellis et al., 2009; Pellis & Bell, 2011; Schleidt & Crawley, 1980).

References

- Adamo, S. A., & Hoy, R. R. (1995). Agonistic behaviour in male and female field crickets, *Gryllus bimaculatus*, and how behavioural context influences its expression. *Animal Behaviour*, *49*, 1491-1501.
- Albert, D. (2011). What's on the mind of a jellyfish? A review of the biobehavioural observations on *Aurelia* sp. jellyfish. *Neuroscience and Biobehavioural Reviews*, *35*, 474-482.
- Armitage, C. J., & Conner, M. (2001). Efficacy of a theory of planned behavior: A meta-analytic review. *Social Psychology*, *40*, 471-499.
- Ashby, W. R. (1952). *Design for a brain*. London: Chapman & Hall Ltd.
- Ashby, W. R. (1956). *An introduction to cybernetics*. London: Chapman & Hall Ltd.
- Ashby, W. R. (1958). Requisite variety and its implications for the control of complex systems. *Cybernetica*, *1*, 83-99.
- Ayala, F. J., & Rzhetsky, A. (1998). Origin of the metazoan phyla: Molecular clock confirms paleontological estimates. *Proceedings of the National Academy of Sciences of the United States of America*, *95*, 606-611.
- Barrett, L. (2011). *Beyond the brain: How body and environment shape animal and human minds*. Princeton NJ: Princeton University Press.
- Beck, J. M., Ma, W. J., Pitkow, X., Latham, P. E., & Pouget, A. (2012). Not noisy, just wrong: The role of suboptimal inference in behavioral variability. *Neuron*, *74*, 30-39.
- Begall, S., Cervený, J., Neef, J., Vojtech, O., & Burda, H. (2008). Magnetic alignment in grazing and resting cattle and deer. *Proceedings of the National Academy of Sciences of the United States of America*, *105*, 13451-13455.
- Bell, H. C. (submitted). Behavioral variability in the service of constancy. *Interna-*

tional Journal of Comparative Psychology.

- Bell, H. C., Judge, K. A., Johnson, E. A., Cade, W. H., & Pellis, S. M. (2012). How is a cricket like a rat? Insights from the application of cybernetics to food protective behaviour. *Animal Behaviour*, *84*, 843-851.
- Bell, H. C., & Pellis, S. M. (2011). A cybernetic perspective on food protection in rats: Simple rules can generate complex and adaptable behaviour. *Animal Behaviour*, *82*, 659-666.
- Bell, H. C., Pellis, S. M., & Kolb, B. (2010). Juvenile peer play experience and the development of the orbitofrontal and medial prefrontal cortices. *Behavioural Brain Research*, *207*, 7-13.
- Bell, W. J. (1991). *Searching behaviour: The behavioural ecology of finding resources*. Cambridge, UK: University Press.
- Bermejo, R., Harvey, M., Gao, P., & Zeigler, H. P. (1996). Conditioned whisking in the rat. *Somatosensory and Motor Research*, *13*, 225-233.
- Bernardet, U., Bermudez i Badia, S., & Verschure, P. F. M. (2012). A model for the neuronal substrate of dead reckoning in arthropods: A comparative computational and behavioral study. *Theory in Biosciences*, *127*, 163-175.
- Bernstein, N. (1967). *The co-ordination and regulation of movements*. Oxford, UK: Pergamon Press.
- Blough, D. S. (1966). The reinforcement of least-frequent interresponse times. *Journal of the Experimental Analysis of Behavior*, *9*, 581-591.
- Böhm, H., Heinzl, H. G., Scharstein, H., & Wendler, G. (1991). The course control system of beetles walking in an air-current field. *Journal of Comparative Physiology A*, *169*, 671-683.
- Bongard, J. C., & Paul, C. (2000). Investigating morphological symmetry and locomotive efficiency using virtual embodied evolution. *From Animals to Animats*,

6, 420-429.

- Breland, K., & Breland, M. (1961). The misbehavior of organisms. *American Psychologist*, *16*, 681-684.
- Brembs, B. (2011). Spontaneous decisions and operant conditioning in fruit flies. *Behavioural Processes*, *87*, 157-164.
- Buzsáki, G. (2006). *Rhythms of the brain*. Oxford, UK: Oxford University Press.
- Carey, T. A. (2006). *The method of levels: How to do psychology without getting in the way*. Hayward, CA: Living Control Systems Publishing.
- Carvell, G. E., & Simons, D. J. (1995). Task-related and subject-related differences in sensorimotor behavior during active touch. *Somatosensory and Motor Research*, *12*, 1-9.
- Clark, D. C., Beshear, D. D., & Moore, A. J. (1995). Role of familiarity in structuring male-male social interactions in the cockroach *gramphadorhina portentosa* (Dictyoptera, Blaberidae). *Annals of the Entomological Society of America*, *88*, 554-561.
- Clark, D. C., & Moore, A. J. (1994). Social interactions and aggression among male Madagascar Hissing Cockroaches (*Gramphadorhina portentosa*) in groups (Dictyoptera, Blaberidae). *Journal of Insect Behavior*, *7*, 199-215.
- Cziko, G. (2000). *The things we do: Using the lessons of Bernard and Darwin to understand the what, how and why of our behavior*. Cambridge, MA: MIT Press.
- Dial, K. P., Jackson, B. E., & Segre, P. (2008). A fundamental avian wingstroke provides a new perspective on the evolution of flight. *Nature*, *451*, 986-989.
- Dickinson, A., & Balleine, B. W. (2000). Causal cognition and goal-directed action. In C. Heyes & L. Huber (Eds.), *The evolution of cognition* (p. 185-204). Cambridge MA: MIT Press.

- Driver, P. M., & Humphries, D. A. (1988). *Protean behaviour: The biology of unpredictability*. Oxford, UK: Clarendon Press.
- Eshkol, N., & Wachman, A. (1958). *Movement notation*. London: Weidenfeld and Nicholson.
- Fend, M., Bovet, S., & Pfeifer, R. (2006). On the influence of morphology of tactile sensors for behavior and control. *Robotics and Autonomous Systems*, *54*, 686-695.
- Field, E. F., Watson, N. V., Whishaw, I. Q., & Pellis, S. M. (2005). A masculinized skeletomusculature is not necessary for male-typical patterns of food-protective movement. *Hormones and Behavior*, *47*, 49-55.
- Field, E. F., Whishaw, I. Q., Forgie, M. L., & Pellis, S. M. (2004). Neonatal and pubertal, but not adult, ovarian steroids are necessary for the development of female-typical patterns of dodging to protect a food item. *Behavioral Neuroscience*, *118*, 1293-1304.
- Field, E. F., Whishaw, I. Q., & Pellis, S. M. (1996). A kinematic analysis of evasive dodging movements used during food protection in the rat (*Rattus norvegicus*): Evidence for sex differences in movement. *Journal of Comparative Psychology*, *110*, 298-306.
- Field, E. F., Whishaw, I. Q., & Pellis, S. M. (1997a). A kinematic analysis of sex-typical movement patterns used during evasive dodging to protect a food item: The role of testicular hormones. *Behavioral Neuroscience*, *11*, 808-815.
- Field, E. F., Whishaw, I. Q., & Pellis, S. M. (1997b). Organization of sex-typical patterns of defense during food protection in the rat: The role of the opponent's sex. *Aggressive Behavior*, *23*, 197-214.
- Finley, J., Ireton, D., Schleidt, W. M., & Thompson, T. A. (1983). A new look at the features of mallard courtship displays. *Animal Behaviour*, *31*, 348-354.

- Friesen, W. O., & Block, G. D. (1984). What is a biological oscillator? *American Journal of Physiology*, *246*, R847-R851.
- Friesen, W. O., Block, G. D., & Hocker, C. G. (1993). Formal approaches to understanding biological oscillators. *Annual Review of Physiology*, *55*, 661-681.
- Golani, I. (1976). Homeostatic motor processes in mammalian interactions: A choreography of display. In P. P. G. Bateson & P. H. Klopfer (Eds.), *Perspectives in ethology* (Vol. 2, p. 69-134). New York, NY: Plenum.
- Grimaldi, B. (2012). *Agent-based vs. robotic simulation: A repeated prisoner's dilemma experiment*. Unpublished Master's dissertation. Università Degli Studi di Torino, Turin, Italy.
- Hebb, D. O. (1949). *The organization of behavior*. New York: Wiley & Sons.
- Himmler, B. T., Stryjek, R., Modinska, K., Derkson, S. M., Pisula, W., & Pellis, S. M. (in press). How domestication modulates play behavior: A comparative analysis between wild rats and a laboratory strain of *Rattus norvegicus*. *Journal of Comparative Psychology*.
- Iwai, N., Sugiura, S., & Chiba, S. (2010). Prey tracking behavior in the invasive terrestrial planarian *Platydemus manokwari* (Platyhelminthes, Tricladida). *Naturwissenschaften*, *97*, 997-1002.
- Johnston, J. (2008). *The allure of machinic life: Cybernetics, artificial life and the new ai*. Cambridge: MIT.
- Jones, G. H., Marsden, C., & Robbins, T. W. (2001). Behavioural rigidity and rule-learning deficits following isolation-rearing in the rat: Neurochemical correlates. *Behavioural Brain Research*, *43*, 35-50.
- Korobkova, E., Emonet, T., Vilar, J. M. G., Shimizu, T. S., & Cluzel, P. (2004). From molecular noise to behavioral variability in a single bacterium. *Nature*, *428*, 574-578.

- Koshland, D. E. J. (1980). *Bacterial chemotaxis as a model behavioral system*. New York: Raven Press.
- Krichmar, J. L., & Wagatsuma, H. (Eds.). (2011). *Neuromorphic and brain-based robots*. Cambridge, UK: Cambridge University Press.
- Lehner, P. N. (1996). *Handbook of ethological methods* (second ed.). Cambridge, UK: Cambridge University Press.
- Lorenz, K. (1981). *The foundations of ethology*. New York: Springer-Verlage.
- Luke, S., Cioffi-Revilla, C., Panait, L., Sullivan, K., & Balan, G. (2005). MASON: A multiagent simulation environment. *Simulation - Transactions of the Society for Modeling and Simulation International*, 81, 517-527.
- Manton, S. M., & Harding, J. P. (1964). Mandibular mechanisms and evolution of arthropods. *Philosophical Transactions of the Royal Society of London Series B - Biological Sciences*, 247, 1-183.
- Marken, R. S. (1986). Perceptual organization of behavior: A hierarchical control model of coordinated action. *Journal of Experimental Psychology*, 12, 267-276.
- Marken, R. S. (1988). The nature of behavior: Control as fact and theory. *Behavioral Science*, 33, 196-206.
- Marken, R. S. (1990). Spreadsheet analysis of a hierarchical control system model of behavior. *Behavior Research Methods Instruments & Computers*, 22, 349-359.
- Marken, R. S. (2009). You say you had a revolution: Methodological foundations of closed-loop psychology. *Review of General Psychology*, 13(2, Sp. Iss. SI), 137-145.
- Marken, R. S., & Mansell, W. (2013). Perceptual control theory as a unifying concept in psychology. *Review of General Psychology*, 17, 190-195.
- Martin, M., Bastien, C., & Albuquerque, P. (2002). Formation of an ant cemetery: Swarm intelligence or statistical accident? *Future Generation Computer*

Systems, 18, 951-959.

- Matsumoto, M., & Nishimura, T. (1998). Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Transactions on Modeling and Computer Simulation*, 8, 3-30.
- May, C. J., & Schank, J. C. (2009). The interaction of body morphology, directional kinematics, and environmental structure in the generation of neonatal rat (*Rattus norvegicus*) locomotor behavior. *Ecological Psychology*, 21, 308-333.
- May, C. J., Schank, J. C., & Joshi, S. (2006). Rat pups and random robots generate similar self-organized and intentional behavior. *Complexity*, 12, 53-66.
- May, C. J., Schank, J. C., & Joshi, S. (2011). Modeling the influence of morphology on the movement ecology of groups of infant rats (*Rattus norvegicus*). *Adaptive Behavior*, 19, 280-291.
- Maye, A., Hsieh, C., Sugihara, G., & Brembs, B. (2007). Order in spontaneous behavior. *PLoS ONE*, 2, e443.
- Mazzoni, A., Broccard, F. D., Garcia-Perez, E., Bonifazi, P., Ruaro, M. E., & Torre, V. (2007). On the dynamics of spontaneous activity in neuronal networks. *PLoS ONE*, 2, e439.
- McBeath, M. K., Schaffer, D. M., & Kaiser, M. K. (1995). How baseball outfielders determine where to run to catch fly balls. *Science*, 268, 569-573.
- McPhail, C., & Tucker, C. W. (1990). Purposive collective action. *The American Behavioral Scientist*, 34, 81-94.
- Mead, C. (1990). Neuromorphic electronic systems. *Proceedings of the IEEE*, 78, 1629-1636.
- Moran, G., Fentress, J. C., & Golani, I. (1981). A description of relational patterns of movement during "ritualized fighting" in wolves. *Animal Behaviour*, 29, 1146-1165.

- Mrosovsky, N. (1990). *Rheostasis: The physiology of change*. Oxford, UK: Oxford University Press.
- Nahodil, P., & Vitku, J. (2012). Learning of autonomous agent in virtual environment. *Proceedings 26th European Conference on Modelling and Simulation ECMS*, 373-379.
- Neuringer, A. (2002). Operant variability: Evidence, functions, and theory. *Psychometric Bulletin & Review*, 9, 672-705.
- Neuringer, A. (2004). Reinforced variability in animals and people: Implications for adaptive action. *American Psychologist*, 59, 891-906.
- Nogueira, F. S., & Lins de Barros, H. G. P. (1995). Study of the motion of magnetotactic bacteria. *European Biophysics Journal*, 24, 13-21.
- Olberg, R. M., Worthington, A. H., & Venator, K. R. (2000). Prey pursuit and interception in dragonflies. *Journal of Comparative Physiology A - Sensory and Neural Behavioral Physiology*, 186, 155-162.
- Oosawa, F., & Nakaoka, Y. (1977). Behavior of micro-organisms as particles with internal state variables. *Journal of Theoretical Biology*, 66, 747-761.
- Park, H., Pontius, W., Guet, C. C., Emonet, T., & Cluzel, P. (2010). Interdependence of behavioral variability and response to small stimuli in bacteria. *Nature*, 468, 819-823.
- Pellis, S. M. (1982). An analysis of courtship and mating in the Cape Barren Goose *Cereopsis novaehollandiae* Latham based on Eshkol-Wachman movement notation. *Bird Behaviour*, 4, 30-41.
- Pellis, S. M., & Bell, H. C. (2011). Closing the circle between perceptions and behavior: A cybernetic view of behavior and its consequences for studying motivation and development. *Developmental Cognitive Neuroscience*, 1, 404-413.

- Pellis, S. M., Blundell, M. A., Bell, H. C., Pellis, V. C., Krakauer, A. H., & Patricelli, G. L. (2013). Drawn into the vortex: The facing-past encounter and combat in lekking male greater sage-grouse (*Centrocercus urophasianus*). *Behaviour*, *150*, 1567-1599.
- Pellis, S. M., Field, E. F., & Whishaw, I. Q. (1999). The development of sex-differentiated defensive motor pattern in rats: A possible role for juvenile experience. *Developmental Psychobiology*, *35*, 156-164.
- Pellis, S. M., Gray, D., & Cade, W. H. (2009). The judder of the cricket: The variance underlying the invariance in behavior. *International Journal of Comparative Psychology*, *22*, 188-205.
- Pellis, S. M., Hastings, E., Takeshi, S., Kamitakahara, H., Komorowska, J., Forgie, M. L., et al. (2006). The effects of orbital frontal cortex damage on the modulation of defensive responses by rats in playful and nonplayful social contexts. *Behavioral Neuroscience*, *120*, 72-84.
- Pellis, S. M., McKenna, M. M., Field, E. F., Pellis, V. C., Prusky, G. T., & Whishaw, I. Q. (1996a). The uses of vision by rats in play fighting and other close quarter social interactions. *Physiology & Behavior*, *59*, 905-913.
- Pellis, S. M., McKenna, M. M., Field, E. F., Pellis, V. C., Prusky, G. T., & Whishaw, I. Q. (1996b). Uses of vision by rats in play fighting and other close-quarter social interactions. *Physiology & Behavior*, *59*(4/5), 905-913.
- Pellis, S. M., & Pellis, V. C. (2006). Play and the development of social engagement: A comparative perspective. In P. J. Marshall & N. A. Fox (Eds.), *The development of social engagement: Neurobiological perspectives* (p. 247-274). Oxford, UK: Oxford University Press.
- Pellis, S. M., Pellis, V. C., Pierce, J. D. J., & Dewsbury, D. A. (1992). Disentangling the contribution of the attacker from that of the defender in the differences

- in the intraspecific fighting of two species of voles. *Aggressive Behavior*, 18, 425-435.
- Pfeifer, R., & Bongard, J. (2007). *How the body shapes the way we think: A new view of intelligence*. Cambridge MA: MIT Press.
- Potegal, M., & Einon, D. (1989). Aggressive behaviors in adult rats deprived of play-fighting as juveniles. *Developmental Psychobiology*, 22, 159-172.
- Powers, W. T. (1973). *Behavior: The control of perception*. London, UK: Wildwood House Ltd.
- Powers, W. T. (2008). *Living control systems III – The Fact of Control*. Montclair, NJ: Benchmark Publications Inc.
- Rescorla, R. A. (1988). Pavlovian conditioning: It's not what you think it is. *American Psychologist*, 43, 151-160.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II* (p. 64-99). New York: Appleton-Century-Crofts.
- Robinson, D. T. (2007). Control systems in Sociology. *Annual Review of Sociology*, 33, 157-174.
- Rosen, R. (1991). *Life itself: A comprehensive inquiry into the nature, origin, and fabrication of life*. USA: Columbia University Press.
- Rosenblueth, A., & Wiener, N. (1945). The role of models in science. *Philosophy of Science*, 12, 316-321.
- Rosenblueth, A., Wiener, N., & Bigelow, J. (1943). Behavior, purpose and teleology. *Philosophy of Science*, 10, 18-24.
- Rothsey, S., & Rohde, K. (2002). The responses of larval copepods and momogeneans to light, gravity and magnetic fields. *Acta Parasitologica*, 47, 167-172.

- Sabarathinam, S., Thamilmaran, K., Borkowski, L., Perlikowski, P., Breski, P., Stefanski, A., et al. (2013). Transient chaos in two coupled, dissipatively perturbed Hamiltonian Duffing oscillators. *Communications in Nonlinear Science and Numerical Simulation*, *18*, 3098-3107.
- Schaffer, D. M., Krauchunas, S. M., Eddy, M., & McBeath, M. K. (2004). How dogs navigate to catch frisbees. *Psychological Science*, *15*, 437-441.
- Schaffer, D. M., & McBeath, M. K. (2002). Baseball outfielders maintain a linear optical trajectory when tracking uncatchable fly balls. *Journal of Experimental Psychology - Human Perception and Performance*, *28*, 335-348.
- Schank, J. C., & Alberts, J. R. (1997). Self-organized huddles of rat pups modeled by simple rules of individual behavior. *Journal of Theoretical Biology*, *189*, 11-25.
- Schank, J. C., May, C. J., & Joshi, S. S. (2004). A biorobotic investigation of Norway rat pups (*Rattus norvegicus*) in an arena. *Adaptive Behavior*, *12*, 161-173.
- Schleidt, W. M., & Crawley, J. N. (1980). Patterns in the behavior of organisms. *Journal of Social and Biological Structures*, *3*, 1-15.
- Schöne, H. (1984). *Spatial orientation: The spatial control of behavior in animals and man*. Princeton, NJ: Princeton University Press.
- Sondergaard, L. V., Herskin, M. S., Ladewig, J., Holm, I. E., & Dagnaes-Hansen, F. (2012). Effect of genetic homogeneity on behavioural variability in an object recognition test in cloned gottingen minipigs. *Applied Animal Behaviour Science*, *141*, 20-24.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Tamburrini, G., & Datteri, E. (2005). Machine experiments and theoretical modelling: From cybernetic methodology to neuro-robotics. *Minds and Machines*, *15*, 334-358.

- Tolman, E. C. (1932). *Purposive behavior in animals and men*. New York: Century.
- Tolman, E. C. (1948). Cognitive maps in rats and men. *The Psychological Review*, *55*, 189-208.
- Uexküll, J. von. (1921). Umwelt und innerwelt der tiere (Environment and inner world of animals). In G. Burghardt (Ed.), *Foundations of comparative ethology*. New York NY: Van Nostrand Reinhold.
- Vácha¹, M., Kvíčalová, M., & Půžová, T. (2009). American cockroaches prefer four geomagnetic positions at rest. *Behaviour*, *147*, 425-440.
- Vancouver, J. B. (2005). The depth of history and explanation as benefit and bane for psychological control theories. *Journal of Applied Psychology*, *90*, 38-52.
- von Holst, E., & Mittelstaedt, H. (1950). Das reafferenzprinzip. Wechselwirkung zwischen zentralnervensystem und peripherie. *Naturewissenschaften*, *37*, 464-474.
- Webb, B. (2000). What does robotics have to offer animal behaviour? *Animal Behaviour*, *60*, 545-558.
- Webb, B. (2008). Using robots to understand animal behavior. *Advances in the Study of Behavior*, *38*, 1-58.
- Wentworth Thompson, D. (1917). *On growth and form*. Cambridge: University Press.
- Whishaw, I. Q. (1988). Food wrenching and dodging - Use of action patterns for the analysis of sensorimotor and social-behavior in the rat. *Journal of Neuroscience Methods*, *24*(2), 169-178.
- Whishaw, I. Q., DuBois, A. T., & Field, E. F. (1998). On the reduction of dodging in mice: A comparison of food wrenching and dodging in rats (*Rattus norvegicus*) and mice (*Mus musculus*). *Journal of Comparative Psychology*, *112*(4), 383-388.

- Whishaw, I. Q., & Gorny, B. P. (1994). Food wrenching and dodging: Eating time estimates influence dodge probability and amplitude. *Aggressive Behavior*, *20*, 35-47.
- Whishaw, I. Q., & Oddie, S. D. (1989). Qualitative and quantitative analyses of hoarding in medial frontal cortex rats using new behavioral paradigm. *Behavioural Brain Research*, *33*, 255-266.
- Whishaw, I. Q., Oddie, S. D., McNamara, R. K., Harris, T. L., & Perry, B. S. (1990). Psychophysical methods for study of sensory-motor behavior using a food-carrying (hoarding) task in rodents. *Journal of Neuroscience Methods*, *32*, 123-133.
- Whishaw, I. Q., & Tomie, J. A. (1988). Food wrenching and dodging - A neuroethological test of cortical and dopaminergic contributions to sensorimotor behavior in the rat. *Behavioral Neuroscience*, *102*(1), 110-&.
- Wiener, N. (1948). *Cybernetics: Or control and communication in the animal and the machine*. Cambridge, MA: MIT Press.
- Wiley, R. H. (1973). Territoriality and non-random mating in sage grouse, *Centrocercus urophasianus*. *Animal Behaviour Monographs*, *6*, 87-169.
- Yang, Z. J., Cameron, K., Lewinger, W., Webb, B., & Alan, M. (2012). Neuromorphic control of stepping pattern generation: A dynamic model with analog circuit implementation. *IEEE Transactions on Neural Networks and Learning Systems*, *23*, 373-384.
- Yin, H. H. (2013). Restoring purpose in behavior. In G. Baldassarre & M. Mirolli (Eds.), *Computational and robotic models of the hierarchical organization of behavior* (p. 319-347). Berlin: Springer-Verlag.
- Young, R. (2000). *Visual control in natural and artificial systems*. Unpublished Doctoral dissertation. University of Surrey, UK.

APPENDIX: WHAT ABOUT OTHER BEHAVIOURS?

In addition to the robbing and dodging system investigated, one other type of behaviour was considered in a small pilot project.

Methods

Subjects

42 male Madagascar hissing cockroaches, born at the university of Lethbridge, were used. The animals were removed from the breeding colony and placed in male-only juvenile colonies once they exhibited sexually dimorphic subgenital plates (the third or fourth instar). Upon reaching the adult molt, indicated by distinctive humps on the pronotum and feathered antennae, males were weighed, measured, and isolated in plastic boxes (21cm long; 314 cm wide; 310 cm high). They were given water and Purina dog chow *ad libitum*, and egg cartons were available for shelter. The animals were kept at 50% humidity level with a 12:12 reverse dark light cycle.

Testing

Subjects were socially isolated for one week prior to the trials to increase the rate of aggression. Subjects were then matched for size and placed together in a housing box (dimensions above) to induce aggressive behaviour. The interior sides of the box were covered in petroleum jelly to prevent escape.

A total of 21 male-male pairings were observed and recorded in the dark using high definition handycams on night vision mode (near-infrared). Interactions were filmed for 15 minutes. In 12 of the trials a handycam was placed directly above the

combatants to observe the basic defensive and offensive maneuvers displayed by each male. The remainders of the trials were observed from this top down angle as well as from the side to gain the angle of each animal in relation to the ground. Because it was believed that the animals were trying to flip each other during encounters, one issue that needed to be addressed was that the substrate upon which the encounters occurred had a low coefficient of friction. To solve this problem, the bottom of the enclosures for the last nine trials was covered with 80 grit sandpaper.

Analysis

Clips of cockroach combat were described and analyzed using a simplified variant of EWMN. The results of the analysis suggested that what each animal is attempting to do is to contact the flank area of the other animal, as this was the area that attackers were most often oriented toward when beginning an attack. This possibility was tested, using different clips, by schematically dividing the cockroachs body into four areas (see Figure 6.1), and then recording where strikes by an attacking animal occurred. It was hypothesized that the reason the animals were targeting the flanks, if this were indeed the case, was because it was the area on the body that is most likely to result in a flip when butted by an attacker. Therefore, in addition to noting where attackers contacted defenders, whether or not the contact resulted in a flip was also recorded.

Results

An ANOVA was performed on the frequency data for contact by body area. The results indicated that the flank area was the preferred target for contact, $F_{3,69} = 25.90, p < 0.05$ (see Figure 6.2). With respect to flips achieved by body area, an



Figure 6.1: Four division of the cockroach body

ANOVA performed on the proportion of flips with respect to contacts per area revealed that, when contacted by an attacker, the flank area resulted in the highest proportion of flips, $F_{3,69} = 3.36, p < 0.05$ (see Figure 6.3).

Discussion

In terms of cybernetics, the cockroaches are not controlling a parameter in the same way as the rats and crickets were during robbing and dodging. Instead, their compensatory behaviour is directed towards maintaining a particular orientation with respect to the other animal, in addition to minimizing the distance to the target.

As can be seen, the examples of combat in Madagascan hissing cockroaches and sage grouse demonstrate how EWMN can provide insight into what the animals are maintaining as constancies during interactions and so provide clues as to what the cybernetic variables may be that are being controlled. Juxtaposing the movements of the two animals and the inter-animal relationship on an EWMN notated page provides a means by which such constancies can be identified (see Pellis et al., 2013), making the arduous task of seeking what animals are controlling much simpler.

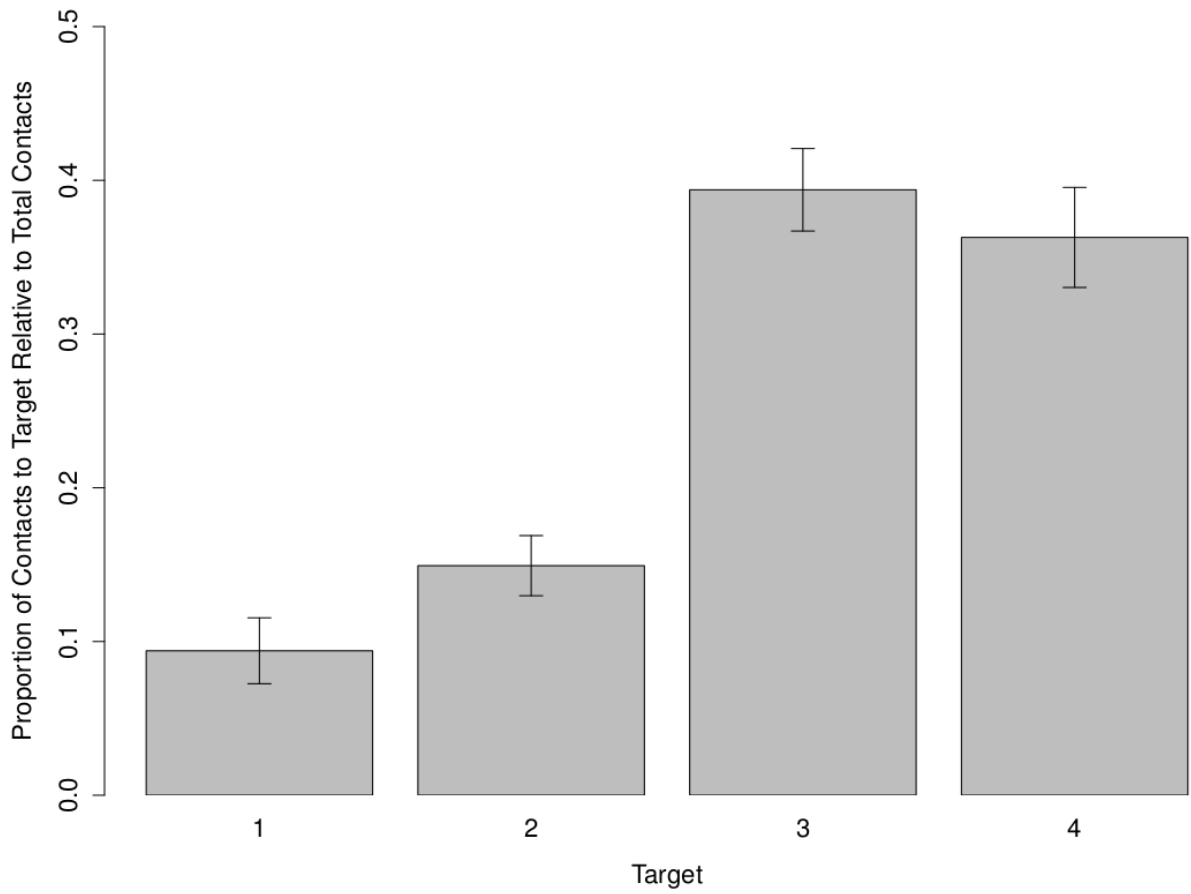


Figure 6.2: Proportion of contacts to each quarter

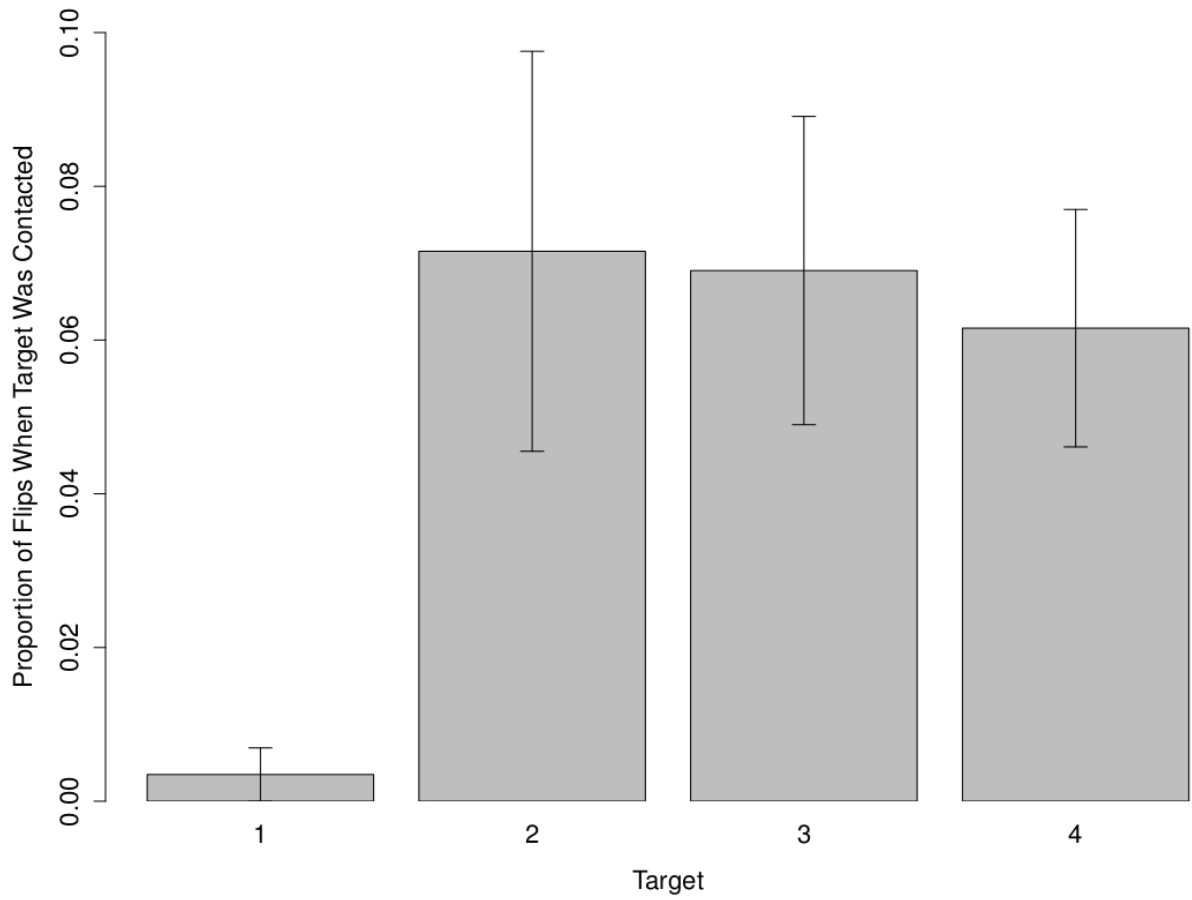


Figure 6.3: Proportion of flips relative to contacts on each quarter