Modelling the Evolution of Warning Signals and Mimicry with Individual-Based Simulations

by

Daniel Wayne Franks

Submitted in accordance with the requirements
for the degree of Doctor of Philosophy.

The University of Leeds
School of Computing

January 2005

Abstract

Individual-based evolutionary simulation models are used alongside mathematical approaches as tools for investigating the selective pressures responsible for the origin and maintenance of animal warning displays and mimicry.

The early chapters of the thesis review the literature on evolution and coevolution in general, as well as setting out a position on the use of computational models in science. The focus then moves to the evolution of warning displays and a review of the biological literature is presented, followed by an evolutionary simulation model. Bright warning displays of defended animals (those possessing a defence such as a sting or toxin) are among the most salient signalling systems in nature. Examples include black and yellow striped stinging wasps, bitter-tasting ladybird beetles, and brightly coloured, toxic butterflies: these conspicuous signals are thought to have evolved as a ‘keep away’ warning to predators. A novel coevolutionary model (Sherratt, 2002a) regarding the evolution of warning displays is replicated and critiqued in detail. The results highlight the utility of a coevolutionary approach but also demonstrate that the theoretical account of the evolution of warning signals remains incomplete.

The thesis then addresses open theoretical issues surrounding the evolution of mimicry (i.e., the adaptive resemblance of one species to another). A review of the biological literature is followed by an evolutionary simulation model. Mimicry is conventionally divided into two categories: Batesian and Müllerian. Batesian mimicry occurs when a species without a defence has evolved an appearance that is similar to that of a defended species. Müllerian mimicry occurs when two defended species evolve to mimic each other. Müllerian mimicries between multiple species are termed *mimicry rings*. The simulation is used to examine the evolution of such rings: first the model is used to look at whether multiple mimicry rings should be expected to form (as in nature), and second it is used to look at the effect of Batesian mimics on the evolution of mimicry rings. The simulation results show that mimicry rings can form because intermediate mutational forms (mutant prey that do not share the appearance of any ring) are at a selective disadvantage, and that Batesian mimics influence the evolution of Müllerian mimicry rings by encouraging those rings to converge.

Finally, theories of warning signals and mimicry are brought together in a mathematical model that looks at whether animals with a defence might evolve traits that are difficult for undefended animals to exploit. This is the first model to implement both mimicry and conspicuous warning displays, and allows us to ask whether defended prey might have evolved conspicuous warning signals in order to shake off parasitic Batesian mimics. The results show that warning signals can, in theory, evolve in response to mimicry. The thesis concludes with a discussion of ways in which individual-based models could be used to take the investigation forward.

Acknowledgements

First and foremost I would like to give my thanks to my supervisor Jason Noble. Jason has kindly given up many more hours than is required of him in order to contribute to my research and significantly develop my academic and professional skills. Further, he has been supportive and enouraging throughout—thanks Jason.

For love, support, and encouragement I would like to thank my fiancée Lindsey Mitchell, and my parents. This thesis would simply not have been possible without you.

Tom Sherratt has been highly influential throughout the latter stages of my PhD. Tom was a good sport about chapter 4, (on which he provided much feedback), a great collaborator for chapter 8, and an invaluable source of feedback overall.

Many more people have been influential to my professional development during the course of the PhD. I would like to sincerely thank the following people: Chris Buckley, Dave Harris, David Hogg, Peter Jimack, Martin Thompson, John Turner, Hywel Williams, all members of the Biosystems group, the administration and computer support staff in the School of Computing, six anonymous referees for the Royal Society of London, and the many people who contributed valuable feedback when I have presented my work at conferences.

To my family.

Declarations

Some parts of the work presented in this thesis have been published (or submitted for publication) in the following articles:

T. N. Sherratt & D. W. Franks. (Submitted). Do unprofitable prey evolve traits that profitable prey find hard to exploit? *Proceedings of the Royal Society of London B: Biological Sciences*.

D. W. Franks & J. Noble. (2004). Warning Signals and Predator-Prey Coevolution. *Proceedings of the Royal Society of London B: Biological Sciences*, vol. 271(1550), pp. 1859-1866.

D. W. Franks & J. Noble. (2003). Batesian Mimics Influence Mimicry Ring Evolution. *The Proceedings of the Royal Society of London B: Biological Sciences*, vol. 271, pp. 191-196.

D. W. Franks & J. Noble. (2003). The Origins of Mimicry Rings. *Artificial Life VIII: The 8th International Conference on the Simulation and Synthesis of Living Systems*, Sydney, Australia.

D. W. Franks & J. Noble. (2002). Conditions for the Evolution of Mimicry. *Poster presented at: From Animals to Animats 7: Proceedings of the Seventh International Conference on the Simulation of Adaptive Behavior*, Edinburgh, UK.

Academia is a collaborative enterprise, and several chapters in this thesis draw on previously published work that I have co-authored with others. For those papers on which my supervisor, Jason Noble, is a co-author, the distribution of work was as follows: identification of need for models (DWF), discussion of ideas and planning of models (DWF & JN), development of simulation models (DWF), primary write-up of paper (DWF), editing and refinement of paper (DWF & JN).

Chapter 8 is based on work presented in Sherratt and Franks (*submitted*). The work distribution for this paper (verified by T. N. Sherratt) was as follows: independent identification of the need for models combining warning signals and mimicry (TNS & DWF), development of the model (TNS with significant input from DWF), implementation of the model (TNS), primary authorship of paper (TNS), further development of manuscript (DWF & TNS).

My collaboration with Noble and with Sherratt has been valuable in developing the ideas behind the work presented here. However, in all cases, the discussions of the ideas and the actual text presented here are my own.

Contents

List of Figures

Chapter 1

Introduction

A warning signal is a conspicuous display made by a species advertising the fact that it has some kind of defence, such as a sting with which to attack a predator, or a toxin that ensures the animal will be unpleasant to eat. Warning signals are a common occurrence in nature, and have set the scene for the evolution of mimicry: the adaptive resemblance of one species to another. One type of mimicry, for example, occurs when a species that does not possess a defence shares the appearance of a species that does possess a defence. There is much that we do not fully understand about the existence of these two phenomena. For example, what benefit would be gained by the first ever warningly coloured individual? In terms of the mimicry relationship outlined above, why does the model tolerate the mimic? The aim of this thesis is to contribute to an answer to these, and more, questions, and to increase our understanding of warning signals and mimicry through the construction of computer models. Before we consider specific theoretical issues regarding mimicry and warning signals, we first need to review the general framework of evolutionary thought upon which the thesis will rely.

1.1 Evolution by Natural Selection

For as long as people have been speculating about the nature of the world, they have noticed that there is a profusion of variety among living things. From blue whales to oak trees, from fleas to flu viruses, organisms show many different ways of surviving in the world. In addition to the enormous variety of forms and behaviours exhibited in nature, organisms also show deep similarities and apparent connections, e.g., the popularity of the quadrupedal body plan, or the presence of a heart for pumping blood.

In an attempt to explain both the profuse variety and underlying similarities among organisms, the Ancient Greek philosopher Anaximander and the Roman philosopher Lucretius first introduced the general concept that all living things are related and that they change over time (Osborn, 1929). In particular, Anaximander's text *On Nature* anticipated the theory of evolution, stating that life started as slime in the oceans and eventually moved to drier places. Thus in a time where the world was typically looked at as being constant, it was portrayed as dynamic.

Much later, Jean-Baptiste Lamarck (1809) brought the next major development towards a workable evolutionary theory. His theory of evolution suggested that organisms changed over generational time to adapt to their environment. Thus, his work went a long way towards developing the current scientific view, although he is now more commonly associated with his flawed theory of heredity: the theory suggested that acquired traits are inherited (i.e., parents passing on traits acquired during their lifetime). However, the biggest problem with the theory is that it lacked a workable mechanism.

Charles Darwin is more well-known, more recent, and ultimately a more radical proponent of the same idea. Darwin sailed as a naturalist on board the Beagle in order to collect, observe and document the flora and fauna of the many regions visited. During this voyage Darwin noted both the variety among and general similarities between organisms. Following the Beagle's return he later worked towards developing a theory that would explain these revealing observations. Consequently, the theory of evolution by natural selection was born, providing an overarching explanation for the variety, similarity, and complexity of organisms. This turned out to be one of the most influential theories of all time. Darwin's book *On the origin of species* ignited a scientific revolution. Further, Darwin had placed man within the animal kingdom and delivered the death blow to self-indulgent anthropocentrism.

The most important distinction between Darwin's and Lamarck's theories of evolution was the issue of exactly what was passed on from parent to offspring. Darwin and Lamarck both described a process of *natural selection* (although only Darwin used the term): since not all organisms are equally well adapted to their environment some will survive and reproduce with more success than others. For Darwin, the variation among organisms is innate and is inherited by later generations. Although Darwin could not identify a precise, low-level mechanism for this process of inheritance, history has shown that his hypothesis was basically correct. In contrast, Lamarck's theory included the flawed idea that traits acquired by an animal during its lifetime could be inherited by its offspring. The great strength of Darwin's theory was that it showed how organisms would evolve to become better adapted to their environment, based only on a few simple assumptions.

Although Darwin's theory of evolution by natural selection was a product of the 1830s, he did not publish *On the origin of species* until 1859. For 20 years Darwin had remained silent for fear of persecution from a predominantly Christian community. However, Alfred Russell Wallace independently came up with an almost identical theory of evolution by natural selection and subsequently posted a draft paper to Darwin—the person he thought best able to evaluate it (Bowler, 1944). This pressed Darwin to publish his well-developed 20-year-old theory.

It is interesting (and reassuring) that Wallace and Darwin independently developed such similar theories. Their theories of evolution both had in common the following central points:

1. Organisms *reproduce*.
2. Organisms *vary* in their traits. [1]
3. Organisms that are well adapted to their environment are more likely to survive and reproduce than their lesser counterparts (natural selection).

Darwin and Wallace both reasoned that the logical outcome of these assumptions was that organisms evolved over time to adapt to their environment. Darwin later realised that adaptation is not entirely the full story (although a significant part of it), and became interested in sexual selection (Darwin, 1871). Post-Darwin, evolution by natural selection has, with its rich explanatory power, consequently become the backbone of much of biology.

It is easy to think of evolution as causing a species to become more and more finely tuned in its response to a fixed environmental problem, e.g., a bird's wings becoming more adapted for flight in air, or a polar bear's fur becoming more effective camouflage in a snow-covered landscape. However, more complex situations arise when two or more species develop successive adaptations in response to the problems posed by each other's behaviour—examples include predators and prey, or hosts and parasites. This process is termed *coevolution*, in which interacting species evolve in response to each other, and will be discussed in chapter 2.

---

[1] Variation is introduced to populations by means such as genetic mutations.

1.2 The Problem

Many animals that regularly interact with others, of the same or different species, have evolved attempts to manipulate or influence the behaviour of those others with signalling displays. There are various reasons behind such displays, such as a sexual display in which a male tries to convince a female of his value as a mate, a cooperative signal with which one animal warns another of a predator, or an aggressive signal used by one animal to threaten another in a dispute over territory. Another reason is to give a warning to other animals—a "stay away" signal. This leads us to one of the central topics of this thesis.

As mentioned earlier, variety is profuse throughout the natural world. Some of the most extravagant structures and behaviours in the animal kingdom appear to be used as displays or ornaments. From bright yellow butterflies bearing tiger-like stripes to peacocks flaunting alluring tail feathers, the world is rich with enchanting natural ornaments. Why do these bright and colourful patterns exist? Wallace (1879a) wrote "To the ordinary observer the colours of the various kinds of mollusks, insects, reptiles, birds, and mammals, appear to have no use, and to be distributed pretty much at random." He then explains "[However,] the colours of animals are of the greatest importance to them, and . . . sometimes even their very existence depends upon their peculiar tints."

Both Wallace and Darwin took an interest in these displays, and in particular, in the observation that many prey species which are harmful to predators (due to possessing a defence such as a sting or a toxin) display bright and conspicuous colourations and patterns. In the early days of evolutionary theory, Darwin (1871) and Wallace (1879a) puzzled over the correlation between defences and bright displays. Darwin wrote to Wallace asking "Why are caterpillars artistically coloured?" (Darwin, 1887). Wallace (1871) reasoned that the bright and colourful displays of toxic species must act as warning signals to predators; something along the lines of "don't eat me—I'm nasty". Darwin (1871) compared the link between warning signals and defence to the ". . . principle that poisons are sold in coloured bottles by druggists for the good of man." Thus the basic story that bright colours act as warning signals for predators was in place from very early on in the history of evolutionary thought. However, the complete answer to the puzzle continues to elude evolutionary biologists to this day: in particular, it is not clear how a warning signal gets started, in that the first conspicuous mutant will be at an enormous selective disadvantage. More generally, it is unclear how animals displaying a warning signal resolve the trade-off between conspicuous display, which makes them more likely to be detected by predators, and the benefits of being cryptically coloured and thus avoiding the attention of predators in the first place.

Darwin and Wallace were also fascinated by the observations of Henry Walter Bates (1862) that harmless species tended to mimic the bright colourations, patterns, and appearances of toxic species. Darwin (1871, p.182) considered such mimicry and warning displays together:

> . . . it is probable that conspicuous colours are indirectly beneficial to many species, as a warning that they are unpalatable. For in certain other cases, beauty has been gained through the imitation of other beautiful species, which inhabit the same district and enjoy an immunity from attack by being in some way offensive to their enemies; but then we have to account for the beauty of the imitated species.

Mimicry has historical significance as it was used as one of the first major pieces of evidence for Darwinian evolution. Studies of warning displays and mimicry were pursued enthusiastically in the decades following the publication of *On the origin of species*. More recently, as Speed (1999) points out, studies of warning signals and mimicry have helped scientists to shed light on many general evolutionary problems such as kin selection (Fisher, 1930), the genetics behind adaptations (Clarke & Sheppard, 1971; Sheppard, Turner, Brown, Benson, & Singer, 1985), evolutionary history (Turner, 1988), shifting-balance and speciation processes (Mallet & Singer, 1987; Jiggins, McMillan, Neukirchen, & Mallet, 1996), and arms races (Nur, 1970; Turner, 1987; Gavrilets & Hastings, 1998). Such studies have also shed light on whether evolution might proceed by punctuated leaps or gradual steps (Poulton, 1909; Turner, 1983). In the early days mimicry was even invoked in arguments over sexual selection (Darwin, 1871). Thompson (1994), in his review of coevolutionary theory, states that (p. 33) “The evolution of mimicry [has] become a part of all the major arguments on evolutionary processes.”

Although mimicry was a key battleground for early evolutionary thought, in modern biology it has become a relatively specialised area. Chapter 2 presents an argument that the reason for this decline in prominence has to do with the rise of more immediate problems for evolutionary theory and a lack of suitable modelling tools for tackling the complexities of coevolution. It is notable that despite the rapid early progress in the decades after Wallace and Darwin published their theories, there have been relatively few milestones in mimicry research in the last century. In particular, the theoretical relationship between warning signals and mimicry remains unexplored.

1.3 Thesis Aims

This thesis will employ computational simulation techniques to develop models of the evolution of warning signals and mimicry. Of course a computer model is no substitute for an empirical result and thus the thesis will be concerned with theoretical advances and clarifications that must await future empirical tests. However, theoretical advances are needed in order to generate hypotheses worth testing, and to develop the theoretical framework in which warning signals and mimicry are understood.

The results of the computational experiments described here turn out to vindicate some of Wallace's original ideas. Wallace is of course very much in Darwin's shadow in the history of evolutionary thinking; the validation of his thoughts on mimicry helps to show that he was just as insightful an evolutionary thinker as his more celebrated colleague.

This thesis has the following general aims:

1. To describe the use of individual-based simulations for studying theories of warning signals and mimicry.
2. To further understanding of how warning signals evolve.
3. To further understanding of how different types of mimicry interact and coevolve.
4. To further understanding of the relationship between warning signals and mimicry.

1.4 Methodology

Chapter 3 spells out the methodology used in the thesis in detail, but, in general, computational modelling approaches are used as a means to explore biological theories. There is a long history of modelling enterprises in science from armchair models using verbal arguments, through mathematical models employing differential equations, to computational models consisting of complex programs. Equations are excellent mathematical tools which are commonly used for biological modelling; successful applications of such techniques in theoretical biology are too numerous to list. However, due to the need to keep mathematical models tractable, their flexibility is inherently limited. With the development of modern computers, individual-based modelling (explicitly modelling individuals and their interactions in computer programs) has become possible. Such a technique is proving useful for exploring evolutionary theories, and is used in this thesis to explore theories of the evolution of warning signals and mimicry.

1.5 Outline of the Thesis

- Chapter 2 introduces coevolutionary theory in a historical context and gives further details of the mechanisms and dynamics behind coevolution. This helps to set the stage for later discussions of warning signals and mimicry, which are essentially coevolutionary processes.

- Chapter 3 introduces *evolutionary simulation modelling* as a technique for developing evolutionary theories regarding biological systems. The chapter puts evolutionary simulation models in context, describes the scaffolding of such simulations, and addresses methodological issues surrounding their use.

- Chapter 4 offers a critical review of predator-prey interactions and theories and models regarding the evolution of warning signals. In particular, the chapter highlights and compares two seemingly opposing theories, which are addressed in more detail in chapter 8.

- Chapter 5 presents an evolutionary simulation model as a replication of a previously published mathematical model of the coevolution of warning signals and predator strategies. Confusions regarding the original model are cleared up and the exact implications of the model are described. The model is additionally extended to relax some assumptions.

- Chapter 6 offers a critical review of theories and models regarding the evolution of mimicry. The chapter presents open questions in the area, and highlights the lack of milestones currently met using models, before suggesting ways in which theories could be advanced using evolutionary simulation models.

- Chapter 7 presents two novel evolutionary simulation models which address one of the key open questions surrounding the evolution of mimicry. In general, what effect do Batesian mimics (unharmful prey that mimic harmful prey) have on the evolution of Müllerian mimicry (harmful prey mimicking each other)? More specifically, the model examines the question of why multiple mimicry rings (mimicry relationships between many species) exist, and what effect Batesian mimics have on the evolution of such rings.

- Chapter 8 attempts to combine theories of warning signals and mimicry. This chapter presents the first model to incorporate warning signals and mimicry, and offers the first formal treatment of Wallace's theory of warning signals. The chapter also contributes to discussion regarding the apparently opposing theories regarding warning signal evolution. To conclude, the chapter presents many ways in which the model can be taken forward in important ways using evolutionary simulation modelling techniques.

- Chapter 9 summarises the conclusions and novel contributions to scientific knowledge that can be taken from this thesis. Finally, the chapter presents a selection of potential future projects, using computational techniques such as evolutionary simulation modelling, that would help to further develop theories of warning signals and mimicry.

1.6 Summary

Darwin's and Wallace's theory of evolution by natural selection is one of the most influential theories of our time. Its rich explanatory power allows us to generate useful predictive theories about most biological phenomena. Ever since Darwin published the theory, scientists have argued over why many animals—in particular those that possess a defence—flaunt bright and colourful displays. Wallace was the first to suggest that such displays might act as warning signals to predators. This is an idea that has persisted to this day. The observations of Müller, that many warningly coloured species tend to mimic one another, was also of interest to Darwin and Wallace from the outset, and played a significant role in the development of evolutionary theory. This thesis aims to use modern computational techniques to develop hypotheses, and to explore and develop related theories, regarding the evolution of warning signals and mimicry.

Chapter 2

Coevolutionary Theory

In the previous chapter we reviewed the general framework of evolutionary thought on which this thesis relies. This chapter takes the discussion a step further; whereas the previous chapter mainly considered evolutionary adaptation to static environments, this chapter considers evolution in a dynamic environment consisting chiefly of other evolving organisms.

Organisms are embedded in diverse and complex environments. A good deal of the environmental complexity is due to the fact that the organism is tightly entwined in various competitive or cooperative ecological relationships, such as the competition for survival between foxes and rabbits. The resulting evolution of the species involved in these types of interspecific interactions is termed *coevolution*. The term coevolution was introduced to the modern evolutionary biology literature by Ehrlich and Raven (1964) although they used the term somewhat loosely. A clear definition of coevolution was offered by Janzen (1980):

> "Coevolution" may be usefully defined as an evolutionary change in a trait of the individuals in one population in response to a trait of the individuals of a second population, followed by an evolutionary response by the second population to the change in the first.

The very first model of interspecific interactions was presented in a footnote by Müller (1879), to describe a mutualistic mimetic relationship between two species (see chapter 5). Not only was this the first model concerning interspecific interactions, and implicitly coevolution, but the it was the first mathematical treatment of evolution in general. Despite the fact that Müller's equations are a clear case of an interspecific interaction, Mode (1958; some 76 years later) is often cited as presenting the first mathematical treatment of coevolution (see e.g., Thompson, 1994). Nevertheless, it became clear that theories of both evolution and the more complex phenomenon of coevolution could benefit from the construction of mathematical models.

2.1 An Entangled Bank

In the final paragraph of *On the origin of species*, Darwin (1859, p. 459) famously acknowledged the evolutionary importance of interspecific interactions.

> It is interesting to contemplate an entangled bank, clothed with many plants of many kinds, with birds singing on the bushes, with various insects flitting about, and with worms crawling through the damp earth, and to reflect that these elaborately constructed forms, so different from each other, and dependent on each other in so complex a manner, have all been produced by laws acting around us.

Darwin was also aware of the significance of coevolution between the sexes, with particular relevance to human origins (Darwin, 1871). Wallace, having independently arrived at his own theory of evolution, was also aware of the importance of interspecific interactions (although none of his statements had such broad applicability as Darwin's). On a subject highly relevant to this thesis, Wallace (1858) suggested that the cryptic colourations of some animals might have something to do with such interactions.

> Even the peculiar colours of many animals, especially insects, so closely resembling the soil or the leaves or the trunks on which they habitually reside, are explained on the same principle; for though in the course of ages varieties of many tints may have occurred, *yet those races having colours best adapted to concealment from their enemies would inevitably survive the longest* (original italics).

Complex interspecific interactions have undoubtedly had significant impact upon the evolution of most organisms, given that most animals and plants are involved in predator-prey relationships.

2.2 The Rise of Coevolutionary Theory

In 1848 Henry Walter Bates, an entomologist (and die-hard beetle collector), conducted field work in the Amazon, where he collated large collections of animals. Bates was a friend and colleague of Wallace and continued to correspond with him during this time, encouraging the development of his theory of evolution by natural selection. After fourteen years in the Amazon, Bates published a famous paper detailing his observations of colour patterns in butterflies. He observed that non-poisonous butterflies tend to mimic the bright warning colourations of poisonous butterflies. Bates reasoned that the non-poisonous butterflies must have evolved, by natural selection, to resemble the poisonous ones (see chapter 6). Bates (1863) hinted at the board applicability of theories of warning displays and mimicry.

> …on these expanded membranes [i.e., butterfly wings] nature writes, as on a tablet, the story of the modifications of species, so truly do all changes of the organisation register themselves thereon. Moreover, the same colour-patterns of the wings generally show, with great regularity, the degrees of blood-relationship of the species. As the laws of nature must be the same for all beings, the conclusions furnished by this group of insects must be applicable to the whole world.

Bates' work was the first landmark in coevolutionary theory: mimicry was an apparently intuitive example of reciprocal evolutionary change in interacting species. However, although evolution was studied rigorously after Darwin's theory was published, it took longer for coevolution and interspecific interactions to receive the same broad attention.

Eleven years after Bates' discovery, Fritz Müller contributed an extension to mimicry theory. Influenced by his field work, Müller (1879) made an argument to show how two poisonous animals might benefit from sharing a warning colouration. Whereas Bates reported a parasitic type of mimicry (where one species benefits at the expense of at least one another), the type of mimicry reported by Müller appeared to be mutualistic. Thus, mimicry offered two forms of coevolutionary interactions: antagonism and mutualism (see section 2.5). As Turner (1995) put it, "Mimicry [is able] to furnish us with a Darwinian model of coevolution … which will explain to us from elementary principles what occurs when two organisms coevolve: potentially, it is an important system for our general understanding of evolution." Thus, theories of mimicry influenced a rise of interest in coevolutionary theory among biologists.

2.3 The Decline of Coevolutionary Theory

Despite the importance of coevolution to the fundamental questions in biology, it was pushed to the wayside in the early 20th century for a long period of time. Thus, students of evolutionary biology were taught only half the problem of evolutionary adaptation. When interspecific interactions and coevolution was taught, it was done regarding only specific cases, rather than as a general principle (Thompson, 1994). Thus, coevolution was pushed back as a focus for evolutionary theorists. What caused this slump in studies of interspecific interactions in evolution? First, evolutionary biology seemed to lack the appropriate tools for modelling such interactions (although Müller had shown that simple mathematical approaches could help). Second, and probably most influentially, there were bigger fish to fry.

Although coevolution is a fundamentally important principle, there were even more pressing and immediate problems for evolutionary theory. Two biological paradigms—the naturalist/biometry (statistical study of biology) camp and the geneticist (study of hereditary and transmission of genes) camp—were apparently at odds (for more detail see Mayr & Provine, 1980). The stances taken by the two camps seemed incommensurable. Many of the problems were down to communication problems and idiosyncratic terminology. However, the problem plagued evolutionary biology. Further, questions such as "If parents' traits are combined in an offspring, then why have the colours of organisms not evolved into a grey blended mess like randomly mixed paint?" were given as a problem for evolutionary theory. Questions such as this were used as sticks to beat evolutionists.

The pathway was left open for a genetical theory of natural selection. This unified theory was provided by one of the most influential evolutionary theorists: Ronald Aylmer Fisher. In his highly influential book *The Genetical Theory of Natural Selection* Fisher unified Darwinism and Mendelian genetics, reconciling biometry and genetics by showing that genetics was actually reinforcing biometrics (Fisher, 1930). This movement is referred to as the *evolutionary synthesis*. Mayr and Provine (1980) explain that:

> The term "evolutionary synthesis" was introduced by [Huxley (1942)] to designate the general acceptance of two conclusions: gradual evolution can be explained in terms of small genetic changes (mutations) and recombination, and the ordering of this genetic variation by natural selection; and the observed evolutionary phenomena, particularly macroevolutionary processes and speciation, can be explained in a manner that is consistent with the known genetic mechanisms.

The integration of Mendelian genetics into evolutionary theory also answered the question of why everything does not evolve into a grey mess: parents' traits are not blended but are inherited in sections of genes from each parent's chromosome. Thus Darwin's hypothesis that innate characteristics of an organism were somehow passed on to its offspring had finally been paired with a plausible mechanism.

Although work on coevolution in general made little progress in the first half of the twentieth century, mimicry was a notable exception. It carried the flag for coevolution with Fisher (1930) devoting an entire chapter to it. The chapter was the only one to deal with interspecific interactions in detail.

2.4 The Revival of Coevolutionary Theory

After the shock-waves of the evolutionary synthesis had settled, coevolution made its way back into the limelight. Not only were the fundamental pre-synthesis questions out of the way, but new techniques had emerged which were well suited to studies of simple coevolutionary interactions.

Game theory is a mathematical technique that was initially developed for use in economics. In particular, the use of game theory to find Nash equilibria[1] was proving useful to evolutionary thinkers. Due to the parallels between economic change and evolution, it became apparent that game theory could be applied to theoretical studies of coevolutionary phenomena. This application was championed by the highly influential and respected biologist John Maynard Smith (see e.g., Maynard Smith, 1982). The availability of game theory as a tool (and later, fast computers and other techniques) led to an increase in work on these topics.

An important example of the success of the game-theoretic approach was in work on the evolution of communication. In an often-cited paper, Krebs and Dawkins (1984) described interactions between species in terms of *mind-reading and manipulation*. (The perspective introduced by Krebs and Dawkins will prove useful later on in this thesis when we develop a theoretical understanding of warning signals and mimicry.) Krebs and Dawkins maintain that when there is a conflict of interest, receivers (the animals receiving the signal) are under selection pressure to critically assess the behaviour of the other in order to exploit any telltale signs of their intentions, thus extracting useful information. On the other hand signallers are under selection pressure to manipulate the receiver into doing what they want. Krebs and Dawkins (1984, p. 390) thus state that "selection will

---

[1] A set of strategies for a game with the property that no 'player' can benefit by changing his strategy while the other players keep their strategies unchanged.

act simultaneously to increase the power of manipulators and to increase resistance to it". Thus, such signals are typically expected to become increasingly extravagant.

Although interactions between animals with conflicting interests seems to follow from the selfish individual view of evolution, mutualistic interactions also occur. For example, a pack-hunting predator may attempt to recruit a conspecific in order to bring down prey too large for either animal to take on alone. Krebs and Dawkins (1984) call the signals associated with these interactions *conspiratorial whispers*. In contrast with the conflicts of mind-reading and manipulation, animals in this situation allow themselves to be persuaded, i.e., the mind-reading and manipulation involved is in the interest of both parties in the interaction. Thus, the signals between such animals can be cheap and subtle whispers.

Since the revival of interest in coevolution, theoretical progress has also been made in areas such as sexual selection, predator-prey interactions, and mimicry, to name just a handful of areas.

2.5 Coevolutionary Dynamics

2.5.1 Antagonism

Biologist Leigh Van Valen proposed what is now known as the *Red Queen principle*. The name comes from the analogy to the Red Queen chess piece in Lewis Carroll's *Through the Looking Glass*, who explained to Alice "…here, you see, it takes all the running you can do, to keep in the same place." Essentially, the principle explains that for an evolving organism, continuing development is needed simply in order to maintain its fitness relative to the species with which it is coevolving (Van Valen, 1973). The only way predators can compensate for a better defence by the prey (e.g., gazelles running faster) is by developing a better offence (e.g., leopards running faster). In turn, prey need to evolve a better defence to escape predators, which will again be met by an improvement in predator offence; and so the cycle continues. Of course, there is often an escalation limit; the cycle has to stop somewhere due to physiological, genetical or environmental constraints. Thus, we do not see animals running quickly enough to break the sound barrier.

Dawkins and Krebs (1979) termed this type of coevolutionary dynamic—where there is a conflict of interests—an *arms race*; an analogy to the game-like politics of the arming of competitive 'super-powers'. This term is commonly used in biology, but its roots in the military analogy should not be taken too seriously; unlike real arms races, in a coevolutionary arms race an old weapon can regain its advantage (that this can happen in the case of Batesian mimicry will be shown in chapter 6). Thus, the coevolutionary "progress" of an organism provides it with no *net* benefit because its competitors are continually coevolving in response. The red queen principle means that the struggle for existence generally does not get any easier; animals never win, they just get a short respite before the next onslaught (Ridley, 1993). However, in some cases one species might reach an area of the coevolutionary fitness landscape that, for one reason or another, its coevolving rival cannot reach (see chapter 8). Much of the complexity we see today in the animal kingdom is the result of coevolutionary arms races between (and even within) species.

2.5.2 Mutualism

The red queen effect is not ubiquitous in coevolution. As Krebs and Dawkins noted in the case of conspiratorial whispers, some coevolutionary interactions are mutualistic, in which the species involved benefit from complementary interactions. Darwin tackled mutualistic coevolution explicitly in discussing his ideas on how bees and flowers interact. Darwin dedicated his follow-up to the *Origin* to this issue and backed up his ideas with some observations (Darwin, 1877). He imagined the size, length and body form of honeybees evolving to better match the shape and length of the corolla tubes surrounding the pollen so as to more efficiently exploit the resource. Further, he imagined the shape and length of the corolla tubes to be modified to best allow bees to pollinate them (Darwin, 1859, p. 142).

> Thus I can understand how a flower and a bee might slowly become, either simultaneously or one after the other, modified and adapted in the most perfect manner to each other, by the continued preservation of individuals presenting mutual and slightly favourable deviations of structure.

Further examples of mutualistic coevolution include symbiotic relationships such as those between sharks and cleaner fish, or between ruminants and their gut bacteria, as well as less direct relationships such as those between fruit-bearing plants that need to disperse their seeds and fruit-eating animals that want a quick meal.

2.6 Coevolutionary Computation

So far we have only discussed natural evolution. However, evolution has inspired computational algorithms that can be used to solve engineering problems. *Evolutionary computation* refers to several techniques that involve reducing evolution by natural selection to its most basic algorithmic form and implementing the result in a computer program. Examples of evolutionary computation systems include genetic algorithms (see chapter 3) and genetic programming. The techniques are used to evolve solutions to problems; particularly where variables need to be optimised in order to maximize some predefined static fitness criteria. A genetic algorithm, for example, works with populations of potential solutions, and differentially reproduces the solutions relative to their fitness before randomly mutating some parts of the solution in order to form the next generation.

Evolutionary algorithms can only be used to tackle problems for which we can clearly define a static fitness function in advance. However, many problems exist for which it is either difficult or impossible to define a static fitness function. This is especially the case in situations where the problem itself is not well defined. Thus, *co*evolutionary computation is being used to address such problems. Seminal work by Hillis (1990) on the coevolution of list-sorting algorithms and difficult-to-sort lists sparked interest in using coevolutionary computation. However, there are many difficulties that arise with the use of coevolutionary computation, usually involving the failure of the two populations involved to set an adequate challenge for each other. Resolving these difficulties is an active research field (for a summary see Cartlidge, 2004).

Bullock (1997) explains the appeal of evolutionary algorithms:

> The attraction of artificial evolutionary design techniques lies in the undisputed majesty of naturally evolved systems. That the organisms which populate the world—complex, economical, and robust solutions to evolutionary problems—were generated through a process of replication with variation is a supremely encouraging discovery. The development of man-made design processes fashioned in evolution’s likeness offers us the chance to harness and direct some of the creative power of evolution.

Evolutionary algorithms are thus one example of how computer models and biological theory can help us to build computational tools that are effective in coping with complex problems.

Note that this thesis does not seek to make a contribution to the area of coevolutionary optimisation, but rather is firmly directed towards biological modelling and the advancement of evolutionary theory. However, the short history of the field of coevolutionary optimisation has shown that there is great potential for cross-fertilisation between researchers with engineering goals and the work of biological scientists (e.g., Watson & Pollack, 2001; Cartlidge, 2004). Thus it is probable that the results of the simulations of mimicry and warning signals presented in later chapters will be of use to coevolutionary optimisation researchers.

2.7 Summary

Coevolution is the evolutionary change in a trait of the individuals in one population in response to a trait of the individuals of a second population, followed by an evolutionary response by the second population to the change in the first. After the publication of Darwin's theory of evolution, people began to show interest in interspecific interactions. Thus people started to take interest in coevolution, and studied phenomena such as the coevolution of mimicry. However, interest in coevolution died down for a period of time due to more pressing problems that were central to the theory of evolution. After these problems were solved by the modern synthesis (the merging of Mendelian genetics and evolutionary theory) research in the area of coevolution began to pick up again, and has been applied to many areas of biology (particularly with the use of models). Recently, computer scientists have begun to take inspiration from nature. One product of this is the development evolutionary computation, which consists of several techniques that involve reducing evolution by natural selection to its most basic algorithmic form and implementing the result in a computer program. With a primer of evolutionary and coevolutionary theory behind us, we are now able to move on to the specific interests of this thesis.

Chapter 3

Evolutionary Simulation Modelling

Chapter 1 briefly discussed the rise of modelling techniques as scientific tools—from verbal armchair models to complex computer simulations. This chapter describes *evolutionary simulation models* and their role in science. Although this chapter is intended to introduce the fundamental aspects of evolutionary simulation modelling as a technique for theoretical biology, it is not intended as a novel discussion of the character of such models or of the methodological issues surrounding the use of evolutionary simulation models. Other such discussions have been carried out elsewhere numerous times (e.g., Bullock, 1997; Noble, 1998; Di Paolo, 1999).

3.1 How Models Are Used in Science

Science is our attempt to understand the universe. This is done by making empirical observations, proposing structured hypotheses to explain those observations, and testing those hypotheses in valid and reliable ways. Theories are, more or less, well-supported descriptions of how we think things may work. Models are closely linked to theories, and serve as analogs of the real system, process, or object. Some theories will find empirical support and prove useful for many years before eventually being replaced by a more parsimonious or more accurate explanation. However, unsuccessful theories are necessarily more numerous than successful ones. What remains is the current body of scientific knowledge, which is open to modification in light of evidence or more fitting theories.

We are most familiar with highly successful theories, such as Newton's theory of gravity, or Einstein's theory of relativity. However, we need to remember that in practice science is about sorting through many potentially useful theories and discarding the ones that do not fit into the current scientific structure (i.e., empirical evidence and other theories).

How do we get from a theory to an experimental test of that theory? Some theories are simple enough that an empirical test is obvious. If we have a theory that it only rains on the weekend, for example, we can easily test this theory by deploying a rain gauge on various days and recording the results. However, some theories are more complex and require a model as an intermediate entity between theory and experiment. For example, say we have a theory about the housing market which states that more houses will be purchased at periods of high inflation due to the interaction between inflation, interest rates, and other economic variables. It is not necessarily possible to test this theory simply by observing house-buying levels and inflation over a few years. Instead we may need to set up a model, perhaps in mathematical form, that instantiates the proposed relationships between the important variables. Thus, in general, testing a theory may require setting up a model that is implied by that theory. In building the model we are forced to be explicit about the assumptions involved in a theory. The model will then, in its turn, allow us to derive the logical implications of that set of assumptions.

One could ask why the model-building process is necessary. Why not simply consider the theory and derive its implications directly? As noted above, for some simple theories this may be possible. But more often the theory is of sufficient complexity that directly grasping its implications is beyond unaided thought experiments. For example, consider the global behaviour of a complex system such as the transport infrastructure of a city. This system's overall performance is a result of the interactions between components such as cars, buses, and trains.[1] Clearly some type of model will be necessary in order to consider theories about how the system operates. In this way, models allow us to get to grips with the conditional logic at the heart of the scientific method which says "If theory X is correct, then we should expect to see Y happen under these circumstances." We build a model in order to derive the events Y that are in fact predicted by the theory. Armed with this knowledge, we build an appropriate experiment and check to see whether Y in fact happens; if it does not, we have less confidence in theory X. One way of disconfirming theory X is of course by experiment: the predicted results Y—which the model allowed us to predict—simply do not occur. Another way of ruling out theory X would be to show that the model $\bar{X}$ derived from the theory does not lead to the hoped-for conclusions. For example, the verbal statement of the theory suggests that the acceleration of a falling

---

[1] Sometimes termed *emergence* (see e.g., Bhaskar, 1978)

object is proportional to its mass, but the mathematical model $F = Gm_1m_2/d^2$ (when reorganised) shows that the object's mass is in fact irrelevant and that all objects in free-fall in the Earth's gravity will have the same acceleration rate.

There are further reasons why the model-building process is necessary. Expressed verbally, theories can be left open to misinterpretation on account of ambiguous linguistic properties. Building a model of the theory gives us a less ambiguous description of the theory that others can step through. Further, the history of science has shown us many times over that we cannot trust our intuitions as to the outcomes of a set of assumptions. For example, post-Darwin, many people believed that one implication of Darwinian evolution would be that individuals were selected to do what is best for their species. However, as we now know, selection typically works at the level of the individual (or gene) rather than at the species/group level.

Models are thus useful in multiple ways. Essentially they make the implications of a theory explicit. The explicit model helps to suggest an experimental design to test the theory, or it may allow us to show (prior to empirical testing) that a theory is internally inconsistent and cannot be expected to explain the phenomenon of interest because it does not in fact predict its occurrence. In situations where multiple theories exist as potential explanations for observed phenomena, models can be used to explore the implications of the theories and thus, to sometimes help us to understand which theory is the more consistent with the observations.

3.2 Artificial Life

This thesis involves constructing computational models to explore theories about the evolution of warning signals and mimicry. Some of the inspiration for the idea of using computer simulations to understand biological systems comes from work over the last few decades in the field of Artificial Life (Alife). The core idea in Alife is that biological systems are of sufficient complexity that they can only be adequately represented in a computer simulation (as opposed to formalisms such as differential calculus that have traditionally been used in biology to model evolution). Readers unfamiliar with the field of Alife may find it useful to look at one of the many introductory texts available (e.g., Langton, 1989; Levy, 1993; Boden, 1996).

There are two main types of Alife positions: these are distinguished as weak and strong Alife (Boden, 1996). The strong Alife position contends that life is a process which can be abstracted away from any particular medium. In essence, strong Alife attempts to synthesise life rather than simulate it. The strong Alife position has its problems, but a detailed critique is beyond the scope of this thesis. In contrast weak Alife attempts to simulate natural processes in an attempt to understand them. Weak Alife is pretty much in sympathy with the purpose of this thesis, i.e., use computational models to capture theories in biology.

Alife has been a creative field generating many new ideas and methods over its short history. It has been valuable in breaking new ground around the idea that computer models are useful in biology. However, it does have some weak points. The conceptual confusion involved in strong Alife leads to a problem for some Alife work in which the goals of the exercise are not clear. If a simulation model is clearly meant as a way of being explicit about a theory, then fine. If a biological principle is adopted in an engineering effort, that is also unproblematic. Unfortunately, without pointing any fingers, some work in Alife sits between these two goals, and is not fully successful as either science (because the modelling relationship is not made clear) or engineering (because the practical usefulness is debatable). However, that is not to say a model fully intended as weak Alife cannot in hindsight prove useful for engineering applications.

3.3 Evolutionary Simulation Models

The area of evolutionary simulation modelling is a combination of research from parts of weak Alife and parts of ecology where computer programs are used in which individuals are explicitly represented. One theme of this thesis is the advocation of evolutionary simulation models for modelling warning signals and mimicry. This section describes what constitutes an evolutionary simulation model.

What do we mean by a simulation? Bullock (1997) states that:

> A simulation is a model that unfolds over time. Rather than constructing static representations of the process under examination, such as flow charts or equations, and relying on human interpreters to simulate the passage of time, or determine the state of the system at some arbitrary time analytically, the simulation designer captures the dynamics of the original process by specifying dynamic mechanisms which govern how the system changes over time. The character of such a simulation's dynamics is determined experimentally, through allowing the simulation to unfold over time.

Mathematical approaches—such as dynamical equations, game theory, and population genetics—are traditionally used for the modelling of biological evolutionary systems. Such approaches are deterministic and the predicted state of the system at some future time *t* can usually be derived directly from the mathematical formulation, whereas evolutionary simulation models unfold over time and need to be analysed using statistical techniques due to the inclusion of stochastic processes such as mutation.

3.3.1 Using Genetic Algorithms For Scientific Modelling

Our evolutionary simulation models need to explicitly represent (populations of) individuals. This is accomplished using agent-based modelling techniques. But how can we evolve the individuals? Luckily, the field of computational optimisation has been using and developing evolutionarily inspired optimisation algorithms for a long time (Goldberg, 1989). One such algorithm, termed the *genetic algorithm* (GA), is typically combined with an individual-based approach to form the heart of evolutionary simulation models.

In order to describe the type of modelling technique used in this thesis, it is necessary to first explain the central component of the modelling technique. Thus, a description of how GAs operate now follows. This section owes much to Mitchell (1998, chapter 1).

3.3.1.1 Encoding

If we are simulating the evolution of various strategies, such as different agent behaviours or agent morphologies, we need to encode these strategies somehow. This is typically done with an artificial genome in which different genes code for different traits.

There is no such thing as a bias-free encoding; as long as we are abstracting from nature in some respect, we will be adding a bias to our encoding. However, we can work to reduce such biases and justify the use of a given encoding method as an adequate way of abstracting from the system of interest.

There are typically two main approaches to the encoding of artificial genomes. The first approach ignores developmental processes and uses a direct one-to-one mapping of genotype (genetic encoding) to phenotype (realisation of the encoding); each gene directly encodes for a corresponding trait. The second approach includes an abstract developmental process, i.e., the genotype is mapped to the phenotype using some indirect process. Thus, in this second approach a simple genetic encoding can produce a wider range of complex phenotypes. However, because of the non-trivial mapping the system would be more difficult to analyse than a system with a direct mapping. Grammars such as Lindenmayer systems (or L-systems), for example, use a collection of rules (such as replace encoding X with traits XYX, and so forth) that are recursively applied to the encoding to create a phenotype that has a fractal structure when its full development is displayed graphically (Prusinkiewicz & Hanan, 1989).[2]

A specific research direction regarding development or morphogenesis would require a non-trivial genotype to phenotype mapping. However, the simple direct mapping approach is far more common. Given that we are increasingly aware of the causal complexities in the relationship between genes and phenotype, how can we justify the use of such simplistic representations in these models?

The use of such direct mappings is not unique to computational models: we are engaging in what Grafen (1982) calls the phenotypic gambit. Although the complexities of genetics and development constrain the short-term evolution of a species, in the long-run selection will overcome any such constraints and favour the phenotype best adapted to a given niche. Therefore, it is legitimate to build models of the evolution of phenotypes, and ignore genetic complexities. For example, if we were interested in modelling the effects of sexual selection on the length of the peacock's tail, we would represent tail length as a single continuous-valued "gene".[3] This is despite the fact that we know that a real peacock's tail length is governed by complex interactions between genes. The gambit involves assuming that in the long run these genetic complexities are irrelevant and that the important thing is the effect of selection on the phenotypic trait. Hammerstein's (1996) streetcar theory of evolution formally outlined this idea.

The phenotypic gambit, now justified, will be adopted in the simulation work described in this thesis. The models will address the evolution of brightly coloured and mimetic phenotypes. The developmental pathways for the evolution of such phenotypes are both poorly understood and not the focus of investigation in this thesis. Thus, a simple phenotypic approach to model construction seems the best way forward.

3.3.1.2 Mutation

The evolutionary process requires variation within a population. The second law of thermodynamics states that any copying system will suffer occasional errors. Thus, mutated individuals will occasionally be introduced to the population. Although mutations are often deleterious (reduced fitness) they are sometimes neutral (same fitness) or even adaptive (increased fitness). The evolutionary process is fundamentally reliant on mutations as a source of variation.

In a GA, mutation operates on a gene with a given probability when an individual is reproduced. Depending on the genome encoding, the mutation operator alters the value of

---

[2] L-systems are thus useful for generating plant-like structures, and in this case the encoding scheme itself might evolve.

[3] This is known in population genetics as the additive-genetic approach.

a gene using a given function. For example, a mutation operation might just flip a binary gene value, or add a random number taken from a Gaussian distribution to the current value.

If the gene values need to be constrained within a certain boundary, then an operator is needed to prevent mutations to values outside of the boundaries. Bullock (1999) showed that we should be careful when dealing with boundary operators. With a simple evolutionary simulation model, Bullock showed that certain boundary operators carry inherent biases. If, for example, the boundary operator simply truncates any value over upper-bound U back to U, then a bias is created where genomes under neutral selection pressure will be biased towards boundary values (i.e., they will tend to stick to the edges). Thus, reasonable solutions involve using a toroidal range of gene values (i.e., the values loop around), or bouncing values off the boundary walls (e.g., if a value is taken to be $+2$ over the boundary it will be placed -2 below the boundary).

3.3.1.3 Recombination

Recombination is a mechanism for combining parent genotypes in a sexual population in order to produce their joint offspring. For optimisation problems there exist different types of recombination operators. However, biological models that implement sexual reproduction typically use *single point* recombination which works as follows. First, a random point along a parent genome is randomly selected. Second, the new genotype (the offspring) is produced by combining the part of one parent's genome that comes before the crossover point with the part of the other parent's genome that comes after the crossover point.

Recombination, therefore, abstracts and simulates sexual reproduction. Although sexual reproduction might be useful for optimisation problems and for biological models that, for example, look at sex-limited traits, asexual reproduction is usually adequate for most models (and is more parsimonious).

3.3.1.4 Selection

As has been discussed, in nature evolution differentially selects in favour of individuals that are better adapted to their environment than their conspecifics, as the weakest individuals have a lower chance of surviving long enough to reproduce. In individual-based models this selective force is abstracted and can be modelled in various ways. Most selection operators work as follows. First, after a predefined period of time or after a set number of events, the fitness of the individuals is assessed with some predefined fitness function or by providing individuals with costs and benefits during the unfolding of the simulation. Second, individuals are selected to reproduce relative to their relative fitnesses. Specific selection functions are used to select between the individuals according to their fitness. The two most typical methods follows:

- *Fitness proportionate selection:* In Holland's (1975) original genetic algorithm, an individual's chance of reproducing was directly mapped from their fitness. This type of selection is typically implemented using the *roulette wheel* approach. Each individual is assigned a metaphorical slice of a roulette wheel proportionate to their fitness. The roulette wheel is then spun and the individual that the ball lands on is able to reproduce. This process is repeated until a full population of offspring is born.
- *Tournament selection:* Although fitness proportionate selection is an effective method, it requires individual fitness to be bounded or scaled somehow. This might be easy if fitness is defined using an objective function, but in less simplistic scenarios this is not always possible. Tournament selection is one way around this problem. Two individuals are chosen from the population at random. A random number is then generated, and if the number is below some predefined threshold then the fittest individual is selected to reproduce, else the less fit individual is selected. Both are subsequently placed back in the population. For a model that uses this mechanism see chapter 5.

The above mechanisms are often invoked to operate on a population after every descrete generation in the model; the fitness of all individuals is assessed and compared at the same time, and reproduction takes place only once per generation. However, some models us a "steady state" system, in which individuals are assessed and removed from the system during any given generation. Although quantitative fitness measures can be used with steady state models to assess the quality of an individual, they are not required. For example, an artificial predator might choose to attack a prey item, in which case it does not survive to reproduce (see, e.g., chapter 7).

3.4 Methodological Issues

3.4.1 How Can We Do Science With Simulations?

We have already seen how models are currently used in science. However, a core question needs to be addressed before we delve into the models presented in this thesis: How can we do science with simulations? A sceptic might argue that all computer programs, including computer simulations, only ever logically rearrange and spit out what you put into them. Therefore, the sceptic might say, they can never make any discoveries or arrive at any new knowledge.

In order to address this issue we need first to look carefully at how simulation is used in science and engineering, and what (if anything) is the difference between the two fields. Second, we need to look at the processes of scientific knowledge discovery and theory change so we can look at how modelling in general, and simulation modelling in particular, might fit in with more traditional methods such as real-world experiments.

There are two ways in which simulations are typically used—for engineering, or for scientific study. The first will be familiar to most people and involves constructing a detailed model based on well-understood scientific principles. For example, wind-tunnel simulation might be used used to assess a new aircraft design before constructing a new prototype. Facts about air pressure, temperature, the details of how air flows over objects, and so forth, *are all already assumed to be known* and the simulation calculates their net effect, e.g., calculating the drag coefficient of a new wing design. Thus, this type of simulation would typically be used for for making reliable and accurate predictions given reliable starting points. The other use of simulations is in science. This involves investigating situations in which the key assumptions are themselves under question. Accurate prediction is unlikely in these cases. Instead, the goal is to investigate ways in which partially understood principles might fit together to form a complete picture and select between principles.

What do we mean by juggling partially understood principles in order to develop a picture? This question in its turn can only be answered by enriching our concept of how scientific discovery works. A naïve view of science holds that facts are out there waiting to be discovered by experiment, that facts once discovered are immutable, and that it is the job of the scientist to simply go out and dig up some facts using experiments. It is beyond the scope of this chapter to include detailed discussion of the philosophy of science regarding such issues. However, the work of philosophers (such as Popper, Lakatos, Quine, and Kuhn) has collectively demolished this simplistic view of how science works. A more informed view of the process of scientific discovery holds that theories are attempts by scientists to make sense of data—that accepted theories will exist in a web of interconnections, where no part of a theoretical edifice is immune to revision and theories will inevitably guide the process of observation.

With this more sophisticated view of science in mind, the role of modelling and simulation starts to become clearer. So we can now return to the original criticism that models only spit out results from what we put in. Although this is true, the notion that this is not useful is unfounded; the implications of a given set of assumptions are not always obvious, and are sometimes counter-intuitive. This thesis takes the stance of Di Paolo, Noble, and Bullock (2000) who argued that “although simulations can never substitute for empirical data collection, they are valuable tools for re-organising and probing the internal consistency of a theoretical position.” By removing assumptions from a model of a current theory, for example, it can be found that a subset of the assumptions is sufficient to give the desired effect. Thus, models are useful for encouraging parsimony and simple theories that undercut more complex theories.

Considering again our enriched view of scientific investigation, we can see that there is always a role for empirical experiment—otherwise our theoretical construct cannot be linked to anything. However, empirical findings also cannot stand or fall alone. The development of scientific knowledge comes from continuing interactions between models and empirical investigation.

3.4.2 Analytic Models and Simulations

What is the difference between analytic models and computer simulations? Why do we need simulations when we have analytic models? Both of these questions are strongly linked and will be answered together.

Analytic models are explicitly expressed using equations, while simulations exist as computer programs. Taylor and Jefferson (1994) and later Miller (1995) suggested that simulations are more explicit than mathematical models because they are typically explained in terms of individuals and include representations of interactions, whereas mathematical models do not. However, as Bullock (1997) points out, equations can be explicitly included in a publication for the reader to fully understand all of the assumptions being made, whereas simulations cannot. Thus, these models are explicit in different ways.

Mathematical equations need to be kept relatively simple for them to remain tractable. Thus, their assumptions and representations are usually highly abstract. A benefit of this comes in the generality of the results; if the model is abstract then it could be a model of many systems that share those general characteristics. However, this generality might mean that the model fails to capture something important about the specific phenomenon of interest. Thus, a general mathematical model creates a general framework, and can be supported with simulation model that relaxes some of the general assumptions.

Simulation models appear to be easier for some empiricists, who may not be mathematically minded, to understand. Such models typically describe interactions among individuals within a population; thus, the simulation can be easily described by spelling out the details of the interactions between individuals. Thus, empiricists can understand the workings of a well-explained simulation without knowledge of computer programming, whereas mathematical knowledge is typically a requirement for the understanding of (even a well explained) formal model. There are historical examples of how mathematical formulations can confuse empiricists: for example, Ralph Meldola reported to Darwin that the reception of Müller's paper, which he had translated for the Entomological Society, had been lukewarm (despite the presence of Henry Walter Bates who developed the first theory of mimicry). Meldola wrote: "... In fact, I do not think anyone grasped the line of argument through inability to follow the simple algebraic reasoning which Fritz Müller has adopted." (Mallet, 2004a). However, simulation descriptions do not allow us to *fully* understand the process used to derive the results without reference to programming code—in which case the reader would need to understand programming. The point here, however, is that simulation models are easier for empiricists to understand at a higher level, and thus design relevant experiments.

Simulation models also allow us to investigate phenomena which are difficult to model using traditional mathematical techniques. We can use simulation models to implement a phenomenon from the bottom-up; we can thus model the low-level behaviours and interactions of individuals and observe the emergent high-level behaviour. For example, we can simulate flocking behaviour by modelling the individual animals within the population and modifying simple rules of how each animal should act depending on the actions of its neighbours. A simulation model might, then, offer us a promising way to get to grips with theories of high-level behaviours arising from low-level behaviours (Noble, 1998).

Simulation models can be useful in supporting and extending mathematical models by relaxing assumptions that when relaxed would render the mathematical model intractable (or at best extremely complex). Assumptions such as infinite population sizes, random mating, no genetic drift, and so forth, can all be easily relaxed. Further, individual-based models can easily deal with non-linearities, complex interactions, spatial environments, and other such complexities which pose difficulties for mathematical approaches. Noble (1998) gives an example, "... we can look at the effects of space and mobility: the fact that an animal is not static but moves about in space and encounters other animals in a non-random way. In equational models space can be captured to some extent, e.g., by imagining that animals are arranged in some abstract topology, but even this becomes mathematically complicated and cannot be taken very far."

Although we need to be careful of artifacts in our simulations, Noble (1998) points out that artifacts can just as easily be found in other modelling approaches too:

> A game-theoretic model is only as good as the strategies that the author has elected to include. It may well be that strategy A is an ESS [stable equilibrium] when considered with strategies B and C, but would not be one if strategy D were included in the model. The use of [evolutionary simulation modelling] techniques does not make this problem go away, but, especially in the case of such powerfully expressive architectures as artificial neural networks, the problem is eased because very many strategies are accessible to evolution.

Some have claimed that “analytic approaches are certainly doomed” (Bonabeau, 1994). This is not a view advocated in this thesis. Although the limitations of analytic approaches are stressed in parts of the thesis, the intent is by no means to undermine the utility of analytic models for future research. Claims of the superiority of one modelling approach over another are unjustified. Instead, such approaches can exist in harmony and support each other. For example, analytic models can set a general framework which can be well understood, and simulation models can be constructed to relax any constraints placed upon the initial model. Although a combination of mathematical models and simulations is desirable, simulations and mathematical models are still useful on their own.

3.4.3 How Should We Develop and Use Simulation Models?

Two good starting points for modelling are:

- Select a system of interest and develop a simple model of it; or
- Take a previous model and extend/modify it.

If the model is a cross-disciplinary endeavour, then it is desirable to develop collaborations with experts in the domain area. However, empiricists often have the tendency of wanting to include much complexity in models, and see some models as being overly simplistic. It is important to abstract away from problems such that we can understand what is going on in the system, so keeping the initial model simple is essential. Then future work should involve incrementally adding more complexity to the model, and examining the effects of the added complexity.

What should we do with our simulation model once it is built? Di Paolo et al. (2000) distinguish three phases for using simulation models:

1. *Exploratory phase:* After the initial simulation model is built, explore different cases of interest, define relevant observables, record patterns, re-define observables or alter [the] model if necessary.
2. *Experimental phase:* Formulate hypotheses that organise observations, undertake crucial “experiments” to test these hypotheses, [and] explain what goes on in the simulation in [those] terms.
3. *Explanatory:* Relate the organisation of observations to the *theories* about natural phenomena and the hypotheses that motivated the construction of the model in the first place, [and] make explicit the theoretical consequences.

These are sensible suggestions which have been adhered to throughout this thesis. For more detail see Di Paolo et al. (2000).

3.4.4 Evaluating Models

An engineering-type model of a practical prediction problem, such as modelling an aeroplane wing design, is easy to validate and test with real data. However, because the assumptions of scientific models are themselves up for question, validation is not so simple.

How can we evaluate our scientific model models? Essentially, model evaluation all boils done to pragmatism and interactions between theory and empiricism. Models clarify theories and give predictions that can be tested empirically; once empirical tests have been done then theories (and thus models) can be selected between depending on whether their assumptions fit the empirical studies and whether their predictions are vindicated. Thus, the development of models and theories is an ongoing and incremental process. We can be optimistic that ecological models inform empiricists effectively, as indicated by the success of models as predictive tools in weather forecasting, economics, chemistry (in particular drug manufacture) and so forth.

Although predictions are are desirable, they are not an essential product of any model. Because no work stands or falls alone, and is part of a greater body of work (past, present and future; Chalmers, 1999) it does not have to give predictions to be useful. For example, it can provide a framework for future research and extensions (which might then produce predictions), influence the way that problems are looked at, or produce philosophical insights.

3.5 Summary

Models, in verbal or mathamatical form, are useful tools for exploring scientific theories, and have long had great success within science. The advent of fast computers has recently allowed us to move on to more complex forms of modelling. Evolutionary simulations are a modelling technique developed within the field of artificial intelligence, and specifically Alife, where individuals are explicitly represented, and evolved with the use of a genetic algorithm.

Mathematical models and evolutionary simulation models both have their own strengths and weaknesses, and in conclusion, different modelling techniques all have their unique problems and benefits and it is not right to argue that one is better than the other. Both types of models are useful and have their place in science, and using equational models and evolutionary simulation models to support each other might offer the most rigour to theoretical studies.

Chapter 4

Warning Signals, Crypsis, and Predator-Prey Interactions

4.1 Introduction

As we have seen, in their struggle for life animals are locked in coevolutionary battles—with conflicts between predators and prey providing a pertinent exemplar. For example, in response to approaches from predators, gazelles perform stotting (a display involving leaping into the air with all four legs held stiff and straight) as a way of signalling their healthy physical condition, and thus high probability of escape, to predators, who are typically deterred by effective displays (Walther, 1969). Many prey species have evolved a variety of defences, in an attempt to counter the foraging strategies of their respective predators.

In some cases, they have evolved appearances that match their background to some extent, thus concealing them from predators. In other cases they have developed a more aggressive method of self preservation: toxins, stings, spikes, bites, irritants, and many other unpleasant defence mechanisms are used to do harm to any attacking predator. As we saw in chapter 1, it has long been observed that many prey species that possess these types of defences display conspicuous and colourful patterns. Wallace (1889) reasoned that such prey must be using conspicuous displays to advertise their defence to predators. This chapter introduces the concepts of prey defence and crypsis and summarises a selection of theoretical work regarding the evolution of *warning signals*: prey displays conveying the existence of a toxin, sting, or such to predators. The term *warning signal*[1] is used throughout this thesis when referring to conspicuous displays (such as bright colourations) of prey species that are protected by toxins, stings, poisons, and so forth. However, the term is not intended to pre-judge whether or not such displays can be considered true signals.

This chapter is organised as follows: first we take a brief look at how predators work, so that we can understand the selection pressure imposed upon prey. Second, we look at the various means by which prey have evolved to counter predation. Third, we look at animal signalling in general, before moving on to discuss problems, models, and theories regarding the evolution of warning signals.

4.2 Predation

Predation is ubiquitous in the animal kingdom, and because the most effective hunters are the most likely to survive, evolution has designed predators with various strategies for finding a meal. Predators benefit from foraging for meals with as little effort and injury as possible. Thus, predators attempt to maximise the correct balance of food intake while minimizing energy usage and injury (Krebs & Davies, 1997). Different stages of predator foraging are broken down below.

4.2.1 Decision Making

A variety of prey species may exist in a predator's environment; some of which would make a good meal for the predator and some of which would be harmful. The different prey species would not all necessarily look alike; there may be, for example, differences between the colourations, patterns, size, and brightness of some of the different species. Here lies the problem for predators: which prey types should be attacked and consumed, and which should be avoided?

When a predator comes across a prey individual it must decide on its next course of action: should it attack the prey or continue searching? Many predators are born naïve as to which prey types (prey of a certain appearance) to attack and which to ignore. Thus, predators typically make their decisions based on experience. If a predator comes across a potential meal which shares an appearance with a previously toxic meal, then there is an

---

[1] The link between prey defences and conspicuous displays is often termed *aposematism*.

increased chance that the predator will avoid it. The more experiences a predator has with a particular type of prey, the more its attack decision for that prey type will be enforced.[2]

4.2.2 Search and Detection

After some time, when a predator has gained knowledge as to which prey types should be eaten and which should be avoided, it then becomes desirable for it to actively seek out the edible prey types and ignore the harmful ones. Thus, once a predator is well educated it would—with the exception of recognition mistakes and re-sampling—avoid prey types associated with negative past experiences even if the prey type presented itself as an easy target by flaunting a conspicuous colouration. Edible prey types tend to be *cryptic* (concealed somewhat amongst the background); thus predators sometimes benefit from specializing and focusing their attention on hunting for those edible but hard-to-find prey. This presents a problem for predators: how can edible cryptic prey be foraged effectively?

Referring to the predation strategies of birds, Tinbergen (1960, p. 316) suggested that "…a specific searching image is assimilated when the birds have had a number of chance contacts with the species in question." This *search image* hypothesis suggests that predators enhance their search by constructing mental templates of the appearance(s) of edible prey; by scanning their surroundings for a particular abstract shape and colour that matches that edible prey, for example, predators are less likely to find prey that do not match the template, but more likely to find the prey that do. A search image forms after repeated detections of a particular prey type (Pietrewica & Kamil, 1975), and is maintained by further repeated detections (Plaisted & Mackintosh, 1995). Thus, search images should lead to higher levels of predation on common prey forms as predators would detect them at a higher frequency. Signal detection theory has proven to be an effective technique for modelling search images (Staddon & Gendron, 1983; Getty, 1985; Getty, Kamil, & Real, 1987). However, current models typically only model one or two prey species. Thus, these models would benefit from models that allow for multiple species.

The reduced *search rate* hypothesis, put forward by Gendrom and Staddon (1983), has become a well regarded alternative to the search image hypothesis (Guilford & Dawkins, 1987). Although the search rate hypothesis shares many of its predictions with the search image hypothesis, in this case predators are thought to simply take more time to search for particular cryptic prey types when they are common in the environment. Krebs and Davies (1997, p. 52) make clear one of the differences between a search image and search

---

[2] Assuming, for now, that its experience is always positive or always negative with each different prey type.

rate: "…a predator offered two equally cryptic prey types concurrently will choose only one if it forms a search image but will be equally likely to find either if it reduces its search rate." However, theoretical studies have shown that the two hypotheses are not necessarily mutually exclusive and can, in theory, occur simultaneously (Dukas & Ellner, 1993). For an overview of further differences between the two hypotheses see Ruxton, Sherratt, and Speed (2004a). Of course, not all predators necessarily use either of the above hypotheses. Evidence is still needed to show that such mechanisms are at work in different predatory species.

4.2.3 Reinforcement Learning

When a predator consumes a prey item one of two things might happen: either the prey item is harmless and the predator gains a meal, or it has a defence, such as a sting or a toxin, and the predator is harmed in some way.[3] Regardless of whether the experience is positive or negative for the predator, it will learn from the experience and alter its attack disposition for prey types with characteristics similar to those of the prey item which it attacked. Predators generalise their experiences, due to predation strategies (Brower, Alcock, & Brower, 1971) or inherent properties of their sensory systems (Mackintosh, 1985). Thus, they can confuse prey species that are similar in appearance.

Different, but not necessarily mutually exclusive, learning mechanisms could be at work when a predator learns about the profitability of different prey types. The most common mechanisms are discriminative learning and associative learning (see e.g., Shettleworth, 1998).

Associative learning, in terms of the type of interactions we are interested in here, involves learning to associate the appearances of a distinct prey form with a level of defence. For example, predators might learn to associate a bright red prey type with a black stripe with the existence of a prey defence, and a brown prey type with the absence of a defence. Discriminative learning, in the case of foraging predators, involves learning simple rules that help to distinguish profitable from unprofitable prey. For example, a predator might learn that any species with a contrasting black stripe generally signals a prey defence. Discrimination learning involves learning to categorise and 'conceptualise' stimuli (i.e., distinguish between what is the same and what is different).

Both types of learning may be adaptive in different contexts. Associative learning might allow a predator to learn an adaptive response to all prey types in an ecosystem containing only a few different prey types, where rules generated by discriminative learn-

---

[3] But perhaps still gains a meal.

ing might prove too general to be as comprehensive. However, in an ecosystem comprising many different prey types a predator may not have the mental capacity to learn a response to each individual prey type. In this situation discriminative learning would allow a predator to learn general rules for how to deal with different prey types, and eventually use associative learning to deal with those prey types that do not fit the rules (Beatty, Beirinckx, & Sherratt, 2004).

4.3 Prey Defences

Given the predation strategies of predators, what has evolution done to increase the chance of survival of edible prey? Typically, edible prey have evolved one of two ways to enhance their survival in the face of danger from predators. These types of defence are outlined below.

4.3.1 Primary Defence

Attempting to avoid encounters with predators in the first place is one strategy that has commonly evolved in prey species to protect them from predation. To do so, they have evolved a variety of means with which to hamper a predator's ability to detect them. Such means include: disruptive colourations (as formalised by, Cott, 1940), where prey appearances have evolved to confuse predators by making the detection of their edges and boundaries difficult; transparency and silvering, and (most controversially) countershading (Ruxton et al., 2004a). However, in many instances increased predator avoidance has been achieved through the evolution of body colours that match those of the surrounding background, providing an effective camouflage (Cott, 1940). This blending of an animal's appearance with its background is termed *crypsis* (Endler, 1988)—the type of primary defence with which this thesis is concerned. Endler (1978) provided the first definition of crypsis, which began "…a colour pattern is cryptic if it resembles a random sample from the background." By reducing the chance of being discovered, cryptic prey can avoid many encounters with predators and thus reduce the risk of being killed or injured in an attack (for evidence of the survival value of crypsis see, e.g., Feltmate and Williams, 1989; Johannesson and Ekendahl, 2002). One case of crypsis has provided one of the most notorious examples of natural selection in action. The industrial revolution caused black forms of the peppered moth (*Biston betularia*) to increase in the population, while the white forms depleted. The industrial revolution caused trees to become dark and sooty, and thus the black moths were more cryptic than the white moths, and were protected to some extend from bird predation.[4] Crypsis is an example of a *primary defence*, in which an attempt is made by prey to avoid any initial contact with predators at the search and detection stage.

4.3.2 Secondary Defence

An alternative strategy adopted by many prey, as an alternative to concealing themselves, is the use of stings, toxins, spikes, and other such traits that would harm an attacking predator. Peruvian poison frogs, *Dendrobates imitator* (Symula, Schulte, & Summers, 2001) for example, have evolved a deadly toxin which would injure any predator that might attempt to make a meal of them. This is an example of a *secondary defence*, in which an attempt is made to deter or deflect predators at the attack stage. Examples of secondary defences include, toxins, quick escape[5], erratic flight, breakable wings, irritants, spines, bad tastes or smells, and stings. Prey with this form of defence (or a subset of it) are commonly referred to as defended, unprofitable, distasteful or unpalatable in the warning signal and mimicry literature. In this thesis, prey that possess some form of secondary defence will be referred to as *defended*, and prey without a secondary defence will be referred to as *undefended*. This definition is used in order to accommodate all forms of anti-predator secondary defence.

The evolution of secondary defences is a difficult and interesting problem (Fisher, 1930) but will not be considered in this thesis. It is an important problem, but its detailed inclusion in the discussion within this thesis would complicate matters prematurely.

4.4 Animal Signalling

In nature, many varieties of signalling occur, from male birds using a complex repertoire of songs to signal their reproductive quality to females (Searcy, 1992), through to vervet monkeys producing various alarm calls to alert conspecifics to predator sightings (Struhsaker, 1967; Cheney & Seyfarth, 1982, 1990), and animals making threat displays in contests over resources such as food, to name but a few instances. For example, Mantis shrimp *Gonodactylus bredini* fiercely contest the ownership of cavities in their coral reefs.

---

[4] There has been much controversy in recent years over whether the selection pressure for this has come from bird populations, but (Mallet, 2004b) points out that none of these refutations are based on new evidence, and that all the existing evidence (see Cook, 2000) points to the peppered moth story being convincing.

[5] Prey that can quickly escape could, in some cases, bring no net reward predators that spend time chasing them, even if they are eventually caught and consumed (Ruxton, Speed, & Sherratt, 2004b).

These disputes can escalate into physical confrontations, but are often resolved with signalling competitions involving claw-spreading threat displays (Adams & Caldwell, 1990). These coevolutionary interactions between signallers and receivers have sparked much interest in the area of evolutionary biology, as they may have been the starting point for the evolution of communication and language (Noble, 1998).

4.4.1 Why Are Signals Reliable?

Why are animal signals reliable? That is, why do most individuals engaged in signalling display honestly, when they can bluff? Tail length in male peacocks, for example, is a reliable signal of the quality of the male. Thus, male peacocks are sexually selected disproportionately by females with a preference for the longer tails. Why then do poor-quality males not bluff the signal and grow a long tail despite their low quality? In turn, why would the females then trust in the signalling system? For a signalling system to exist it needs to maintain its reliability, or else receivers would simply pay it no attention. The literature regarding this problem is vast, but some general solutions are discussed below. For more details on the contexts in which minimal cost signalling is stable see, for example, Maynard Smith and Harper (2003).

4.4.2 The Handicap Principle

Zahavi (1975, 1977) hypothesised that the reliability of a signalling system could be maintained if the signals were costly to produce. Consider the following example: natural selection acts upon female peacocks to select high quality mates from amongst groups of males of various qualities. However, the females have no way of *directly* assessing the quality of males. Instead, males advertise quality with flamboyant plumage and a long tail. This display is likely to be costly (in terms of survivorship) to the male: it may be energetically costly, make male peacocks easier to catch, make them bright and thus easy for predators to see, or a combination of the above. We know that in nature the length of male peacocks' tails are correlated with their overall genetic quality—but how do the costs associated with tail length lead to this correlation? Zahavi theoried that lower quality signallers would not be able to afford the costs that higher quality signallers would be able to bear. The original formulation of the handicap principle sparked much controversy, as it was still difficult to see how Zahavi's assumptions could lead to stable honest signalling.

An influx of models followed Zahavi (1975) in an attempt to shed light on the controversy. Many of these models appeared to show that the handicap principle did not enforce honest signalling (for a discussion see, Noble, 1998). However, a landmark paper emerged from all of this to inject life back into the theory. Grafen (1990) used game theory to model costly signalling and found that the handicap principle could work under the following conditions: 1) Signals must bear a cost, 2) The fitness cost of any signal must be greater for higher quality signallers, 3) The fitness benefit of a positive repsonse must not be greater for lower quality signallers. In terms of the example, the fitness cost of extending one's tail by an extra centimetre must be higher for unhealthy or weak males than for healthy strong ones. Thus, high quality male peacocks can give costly displays that improve their chances of reproduction with minimal cost to their survival, whereas low quality males that adopt the costly trait would suffer a bigger blow to their chance of survival. Consequently, males with a long tail and bright plumage are more likely to be high quality, and so females should trust the signal. Grafen found a way for the handicap principle to work despite the potential advantages to individuals that lie. Signalling costs can indeed enforce honesty and reliability.[6]

4.4.3 Indices of Quality

Another reason that a signalling system might remain reliable is that a signal might be inherently difficult or impossible to fake. Deep roaring in red deer, for example, might count as an unfakeable signal of size and strength as the low pitch of the roar is a physical result of the size of the deer (Guilford & Dawkins, 1991). Thus the signal can be used as a reliable index of the quality being signalled (Maynard Smith & Harper, 2003).

4.4.4 Minimal-Cost Signalling

Maynard Smith and Harper (2003, p. 37) describe minimal-cost signals as "…a signal whose reliability does not depend on its cost (i.e., not a handicap), and which could be given by any signaller (i.e., not an index)." Maynard Smith and Harper (2003) point out four contexts in which minimal-cost signalling can be evolutionarily stable:

1. Signaller and receiver place the possible outcomes of the interaction in the same rank order.
2. Dishonest signals are punished (Maynard Smith & Harper, 1988)—in other words, there is a heavier cost to dishonest signals than to honest ones. Thus, the interest here is to find mechanisms for punishing dishonest individuals (see chapter 8).

---

[6] There are three main categories of handicaps (pure epistasis handicap, conditional handicap and revealing handicap), which will not be detailed here. For more information see, for example, Noble (1998).

3. Common interests exist, such as when competing individuals are likely to be related, or when both parties benefit from resolving the competition without it escalating into a fight.

4. Repeated interactions take place and receivers can recognise and remember individuals (Silk, Kaldor, & Boyd, 2000).

4.4.5 Conventional Signalling

Although theories, such as the handicap principle, explain why costly signals are reliable, ‘conventional signals’ that are not costly can be reliable too (e.g., Enquist, 1985; Noble, 2000). Hurd and Enquist (1998) describe this type of signalling: “Conventional signals are those for which the meaning and form of the signals are associated by arbitrary convention, it is communication in its most pure sense.” For example, vervet monkeys make different alarm calls corresponding to the type of predator that they have detected (i.e., there is a leopard call, a snake call, and so forth). Because the calls are only arbitrarily connected with their referents (i.e., they could just have easily used a different sound for each predator), the signals are conventional.

4.5 Warning Signals

The bright warning displays of defended species are amongst the most salient signalling systems in nature, from black and yellow-striped stinging wasps, through to bitter-tasting ladybird beetles, and brightly coloured toxic butterflies (Joron, 2003), warning signals are spread throughout the animal kingdom. As mentioned in Chapter 1, Wallace was the first to hypothesise that these conspicuous displays have evolved as a ‘stay away’ warning to predators. Clearly, such displays are in the mutual interests of defended prey and their potential predators; the prey would benefit from increased survival if the predators avoid them, and the predators would not attack them and suffer the effects of the prey’s defence (such as an injury, an upset stomach, and so forth). The link between defence and conspicuousness referred to as *aposematism* (Poulton, 1890) or a *warning signal*.

Warning signals might be even more prevalent in nature than it first appears, because different warning signals have evolved in response to the particular nervous systems of predators. To give an analogous example, Heiling, Herberstein, and Chittka (2003) argue that the crab spider (*Thomisus spectabilis*) is cryptic to humans but highly conspicuous to the bees on which it preys. This bright display is thought to increase the attractiveness of the flower on which the spider rests, thus luring unwitting bees. Although the crab spider's display is not a warning signal, it illustrates the point that warning signals might be more wide spread in nature than is immediately apparent.

There are many different definitions of what constitutes signalling. Warning signals might be categorised as true signals under some definitions, but not under others. However, many books on the subjects use mimicry in their key discussions of signalling (e.g., Maynard Smith & Harper, 2003) and it is clear that, whether or not aposematic species are indeed signalling, theories of the evolution of warning signals and of signalling in general may mutually inform one another (see Chapter 8). Although warning signals are classified as true signals under many definitions, crypsis generally is not (with a few exceptions such as, Hasson, 1994). Crypsis could almost be thought of as an anti-signal.

Conspicuous warning signals can be costly as they make prey easier to detect (due to bright colourations) and catch (if prey are slow moving or aggregated). Thus, on the face of it, warning signals appear to be an example of the handicap principle at work; costs associated with the signalling look to be responsible for its maintenance. Guilford and Dawkins (1993) explored the question "Are warning colours handicaps?" After detailed discussion, which we will not echo here, they conclude that warning signals are not handicaps. Ruxton, Sherratt, and Speed (2004, p. 86) summarise the argument:

> In many cases warning signals incorporate no components of the defence being advertised . . . When displays of defences are decoupled in this way the form of an individual's display does not necessarily indicate anything about its underlying unprofitability [or defence]. For this reason many warning displays can not function as handicap signals (see Guilford and Dawkins, 1993).

A similar argument could also be used to discount warning signals as minimal cost signals. So are warning signals an instance of conventional signalling? It would seem not—warning signals typically make use of conspicuous traits, and not just any arbitrary colour pattern; thus the signals used are unlikely to be arbitrary.

But surely warning signals have to be classified under one of these schemes? Unless warning signals fit clearly into one of the categories, it is probably a mistake to attempt to shoehorn any example of signalling into one of these categories. There are many examples of other signalling systems that do not fit comfortably into these categories (see Noble, 1998). Our current vocabulary for describing signalling systems is limiting, and needs extending. A detailed discussion of what category warning signals can be placed within is beyond the scope of this thesis, and would merit a paper in its own right. Such a discussion would also benefit from a formal model (Guilford & Dawkins, 1993).

Crypsis and warning signals are often discussed together as contrasting and alternative anti-predator strategies; crypsis reduces the frequency of encounters with predators, whereas warning signals increase the frequency of encounters with predators. Thus, they both exist (at varying degrees) at opposite ends of a detectability spectrum. The cost-benefit trade off between crypsis and conspicuousness is, therefore, often the central topic of warning signal theory.

4.6 A Review of Models of Warning Signal Evolution

Despite over a century of interest in the evolution of warning signals, it is clear that the theories still have much room for development. The study of warning signals is one long search for a general explanation of the phenomenon (Härlin & Härlin, 2003),[7], and theories have commonly been contested and expressed with verbal arguments, mathematical models and computer models. Some of the problems addressed are detailed in this section.

4.6.1 The Stages of Warning Signal Evolution

How did warning displays come to signal the presence of a defence in prey? As Turner (1984, p. 1) states "Some species develop a chemical defence and leave it at that." Some species go even further and evolve crypsis, combining both primary and secondary defences. Taking an example from (Fisher, 1930), Turner continues "…the Buff-tip moth *(Phalera bucephala)* is distasteful, but like the majority of palatable species, it is cryptically coloured." It would appear that the Buff-tip moth has taken full advantage of both primary and secondary defence and thus, one might speculate, would enjoy an increased survival advantage as a result. Why, then, have some prey adopted conspicuous displays such as bright and colourful markings, strong smells, slow movement, and aggregations, that are so common in warningly coloured prey? Surely conspicuous prey endure a higher risk of attack from predators? How can this evolve? "Surely it is better not to be seen at all, than to be somewhat mangled before being dropped from the beak as too hot to handle?" (Turner, 1984, p. 1). Ultimately scientists in the area seek to explain the observation that undefended prey tend to be cryptic while defended prey tend to be conspicuous (but see Wüster, 2004).

The question of how warning signals have evolved in defended animals can be split into two stages (Speed, 2001b). However, the two stages could also be considered to-

---

[7] It is, of course, possible that there are many reasons for the evolution of warning signals, some of which would apply in specific cases.

gether. The stages are:

- Stage One: Speed (2001b) referred to this stage as "the problem of the lone mutant". The stage considers the genesis of warning signals. How might a warning signal arise in a population of defended cryptic (or just inconspicuous) prey? This stage is difficult to explain because, in a cryptic population, an initial mutant prey form would be conspicuous (and thus easy to spot) as well as rare. Because there would not be enough of the novel prey type to elicit avoidance in sampling predators (Lindstrom, Alatalo, Lyytinen, & Mappes, 2001), this stage poses a problem.
- Stage Two: This stage considers the spread of conspicuous traits through a defended population. How, once a defended conspicuous mutant survives stage one, does the warning signal spread to fixation within a population, and how is it maintained? This stage has typically received more focus than stage one.

Is warning signal evolution gradual or punctuated? Stage one of warning signal evolution, or 'the problem of the lone mutant' is the scenario most typically considered for the initial evolution of warning signals (where a rare mutant becomes conspicuous by some major mutation and needs to survive to reproduction). The possibility of gradual evolution (e.g., Huynen, Stadler, & Fontana, 1996) is often neglected. As mentioned earlier Yachi and Higashi (1998) found feasible conditions for the *gradual* evolution of warning signals; an idea that is rarely considered in warning signal theory—probably due to the results of some 'novel world' experiments which show that the psychological properties of modern-day predators do not allow for it (Lindström, 1999). But there may still be room for considering the gradual *co*evolution of predators and prey. It has also been suggested that a lone aposematic mutant might be helped to survive because some predators are reluctant to attack novel or unfamiliar prey (Coppinger, 1969, 1970). Of course, it is completely possible for a lone conspicuous mutant to survive to reproduction by chance (Speed, 2001b).

4.6.2 Why Are Warning Signals Conspicuous?

As mentioned previously one of the key questions is why conspicuous traits are used to convey warnings. Ruxton et al. (2004a, p. 87) suggest, "…first because it confers reliability, especially for prey that gain from being freely exposed in their environments; and second because it directs predator attention to the possession of some aversive component."

Conspicuous animals are not necessarily warningly coloured (e.g., they may be conspicuous and undefended). However, it would be difficult for undefended prey to maintain a conspicuous state, as predators would quickly learn to filter them out (see chapter 5); conspicuous undefended prey would suffer higher predation rates than cryptic undefended prey. It is generally agreed that the evolution of a secondary defence therefore precedes the origin of warning signals. Without a defence there is nothing for the prey to signal about (Harvey & Paxton, 1981; Guilford, 1988; Sillén-Tullberg, 1988; Guilford & Dawkins, 1993; Alatalo & Mappes, 1996; Lindström, 1999).

Two camps are currently emerging regarding predator psychology. The contemporary explanation is that defended prey have evolved conspicuous colourations because they are inherently easier for predators to learn (the receiver psychology approach) as a signal of defence (Guilford, 1988; Mallet & Joron, 2000); that the particular psychological biases that might facilitate or support the evolution of warning signals are 'just one of those things'—biased in favour of particular stimuli for reasons that are not necessarily directly connected with signalling. For example, the fact that predators learn to avoid conspicuous prey more quickly than cryptic prey might be the result of a natural predisposition of predatory nervous systems. An alternative explanation considers the possibility that the biases might have coevolved with prey warning signals, and that defended prey have evolved conspicuous colourations to make it easy for predators to distinguish between defended and undefended prey. This leaves a question: "which came first, conspicuousness or special psychological effects?" (Ruxton et al., 2004a). Alternatively, did they gradually coevolve together? These two alternative explanations are detailed below, along with a flavour of the work carried out in each area (also see Chapter 8).

4.6.3 Conspicuousness: Exploiting Receiver Psychology

Many contemporary theories of the evolution of warning signals have stressed the importance of receiver psychology (for reviews see Guilford, 1988, 1990; Endler, 1991; Mallet & Joron, 1999; Speed, 2000). In particular, much theoretical research examining the evolution of warning signals places emphasis on 'special' psychological aspects of predators (Guilford, 1990; Guilford & Dawkins, 1991). The idea is that aspects of receiver psychology, such as biases in learning rates for different prey types, may select for the evolution of warning signals. Relevant aspects of receiver psychology include: phobias, avoidance learning, prey recognition, memory, recognition errors, and so forth. It has, for example, been shown numerous times that most predators learn to avoid conspicuous defended species faster than they learn to avoid cryptic defended species (Gittleman &

Harvey, 1980; Sillén-Tullberg, 1985; Guilford, 1986; Rowe & Guilford, 2000; Roper & Wistow, 1986; Lindstrom et al., 2001).

Associative learning (see section 4.2.3) is the type of learning typically associated with studies of receiver psychology and warning signals. This hypothesis—that warning signals have evolved to take advantage of the psychological predispositions of predators—seems to be the most generally accepted explanation for the evolution of warning signals.

Servedio (2000) used an analytic model and Speed (2001b) used a stochastic computer model to investigate the effects of systematically varying psychological parameters (such as learning and forgetting) on the predation rates for cryptic and conspicuous prey (stage two of warning signal evolution). Servedio's model predicted that warning signals are most likely to evolve when prey are highly defended and learning is instantaneous. Speed's model predicts that the parameters for neophobia, learning, and forgetting rates all contribute to the survival of aposematic prey. However, defended prey typically enjoyed higher a survival rate when conspicuous rather than cryptic, and when predators forget about cryptic prey more quickly than conspicuous prey. This suggests that memory could be an important factor in the evolution of warning signals (for evidence supporting this see, e.g., Roper, 1994). There are currently no models allowing predator educational biases to evolve in response to aposematic prey.

Leimar, Enquist, and Sillén-Tullberg (1986) developed a fairly comprehensive theoretical framework which explored the evolution of conspicuousness and secondary defences in prey. By systematically varying psychological and ecological parameters—such as level of defence, conspicuousness and kinship (or family grouping) in prey (Fisher, 1930), and learning rates and recognition ability in predators—they found the evolutionary stable strategies for the evolution of warning signals. Leimar et al. (1986) found that warning signals would be selected for particularly if the kinship was high and learning rates were higher for conspicuous prey than for cryptic prey. The question of why warning signals use conspicuous traits, rather than some other trait (such as an arbitrary but inconspicuous colour pattern) as the main signal component is addressed in more detail in chapter 8.

4.6.4 Conspicuousness: A Reliable Indicator of Defence

Although the previous theory of the evolution of warning signals is plausible Sherratt and Beatty (2003) point out an important coevolutionary caveat: "Whatever the underlying cause of aposematism, it is likely that predators would evolve an enhanced psychological predisposition to learn to avoid conspicuous prey precisely because such prey tend to be defended (Turner, 1984; Sherratt, 2002a)." In other words, the previous theory may be correct, but the fact that predators may have the biases necessary for defended prey to exploit by becoming conspicuous cannot be taken as evidence for the theory. This is because the biases may have coevolved *after* defended prey developed warning signals. Turner (1984) states that "It remains to be seen which of the mechanisms, innately fast learning of bright colours, innately fast learning of contrasting colours, fast learning of prey that is different, greater initial rate of attack on conspicuous colours, is the significant mechanism in the origin of aposematic colouring [or warning signals]. Although some experiments will be easy to devise. . . the coevolution of the vertebrate nervous system with the insects may have made it impossible for us finally to disentangle the problem." Thus, Turner acknowledged that predator psychology may have coevolved with prey morphology. Although Turner appeared sceptical about resolving the problem, currently emerging approaches and modelling techniques, such as evolutionary simulation models, can be used to steadily build up and develop more complex models of the coevolution of warning signals and predator psychology, in the hope of disentangling the situation.

An emerging alternative to the receiver psychology perspective is that of warning signals as a reliable indicator of a defence of prey. This theory, rather than looking to inherent receiver biases for an explanation, contends that defended prey evolved conspicuous colourations to reinforce the reliability of the warning signal by distinguishing themselves from cryptic undefended prey. Thus, predators should be more likely to believe that a prey item is defended if its warning signal consists of reliable traits (such as conspicuous traits). However, as Sherratt and Beatty (2003) ask: why should selection on defended prey to avoid confusion with undefended prey act upon conspicuousness and not so much on other traits such as pattern or colouration? Because this approach centres around defended prey distinguishing themselves from undefended prey, it emphasises the relationship between mimicry and warning signals. Surprisingly there has been little discussion of the role of mimicry in theories of warning signal evolution.

Overall, this alternative approach contends that warning signals (and in particular conspicuous traits) have been selected for their reliability as indicators of defence, rather than to capitalise upon any inherent psychological biases of predators. Sherratt and Beatty (2003) point out the advantages of looking different from undefended prey: ". . . predators are more likely to learn to react appropriately to a signal if it is a reliable one (Zahavi, 1975; Zahavi & Zahavi, 1997), and being distinguishable is essentially creating a reliable signal." Thus, discrimination learning, rather than just associative learning, is important to this approach.

Although this perspective on the evolution of warning signals has been embraced more recently, it follows from the ideas of early scientists such as Wallace and Fisher. Wallace

(1879b) for example, following a point he hinted at earlier in Wallace (1867), stated that "…eatable insects (if soft and defenceless) are almost always protected by obscure or green tints harmonising with their surroundings. Evidently, therefore, the best way to distinguish the uneatable kinds would be that they should be of gay and brilliant tints, strongly contrasted with their surroundings, and readily distinguishable from a considerable distance." Fisher (1930, p. 148) appeared to share this view, stating that "To be recognised as unpalatable is equivalent to avoiding confusion with palatable species". This point may have more general implications. Referring to animal signals in general Guilford and Dawkins (1991, p. 3) suggested that "Aspects of a signal that lead to its being increasingly detectable may also lead to its being increasingly discriminable, whilst how memorable a signal is may in turn be affected by how discriminable that signal is from other signals."

Sherratt and Beatty (2003) developed an experiment to look at how computer-generated prey items evolved when subject to selection by humans. They note that (p. 384) "…[this] perspective is very different from the contemporary approach that has tended to focus on the evolution of effective educational aids in defended prey alone rather than on selection to maximise the phenotypic difference between defended and undefended prey." In order to test this hypothesis they allowed volunteers to act as predators on virtual prey. Defended and undefended prey species of different levels of conspicuousness existed in the programme, and were presented to the predators. They could then choose whether or not to attack each the prey item, in a manner that maximised a score; volunteers received a reward for attacking an undefended prey item and a deduction for attacking a defended prey item. Surviving prey reproduced and mutated with some probability, thus prey colourations could evolve. The results showed that defended prey consistently adopted conspicuous traits that allowed them to be distinguished from (typically cryptic) undefended prey. Sherratt and Beatty (2003) conclude that defended prey only evolve to become conspicuous in a system that includes unpalatable prey. Thus, they point out that mathematical models of warning signals that exclude undefended prey (e.g., Harvey, Bull, Pemberton, & Paxton, 1982; Sillén-Tullberg & Bryant, 1983; Yachi & Higashi, 1998) may have excluded what is important.

This theory is a viable alternative to the receiver psychology approach, but does it explain why defended prey use conspicuous signals, and not just some trait unrelated to conspicuousness? The idea that warning signals have evolved in defended prey to distinguish them from undefended prey, and the contrast of this theory with the receiver psychology approach, is returned to in greater detail in Chapter 8.

4.6.5 Other Approaches

Many models and theories of warning signals have been developed that do not cleanly fit in with the above categories. A selection of these models is briefly discussed below.

Yachi and Higashi (1998) presented a mathematical model which looks at the evolution of warning signals using *peak-shift* theory (see also Hanson, 1959; Mallet & Singer, 1987; Gamberale & Tullberg, 1996). Peak-shift is a psychological property related to discriminative learning. Experimental trials have shown a displacement in many animals' generalization gradients, meaning that inherent biases in their sensory systems might lead to preferences for signals that are slightly different from the stimuli that they have been trained on. Yachi and Higashi (1998) use a model to show that a peak-shift in predator generalization could serve to provide selection pressure for greater conspicuousness. They conclude that a predator's generalization is arbitrarily shifted in favour of conspicuousness then a *runaway* (self reinforcing) process should expected, where prey evolve a higher level of conspicuousness to take advantage of the predator's generalization peak and the predators' generalization peak shifts to accommodate the new prey appearance. These results have an extra element of interest in that they can explain the *gradual* evolution of warning signals from an initially cryptic population. The possibility of gradual evolution is commonly neglected in theoretical studies of warning signals, and the typical approach is to adopt the assumption of aposematic evolution in two stages (see 4.5) where the genesis of a warning signal is a large mutation of a single individual. There may be alternative unexplored ways in which prey might evolve warning signals by gradual adaptation or runaway processes.

Not all defended prey are conspicuous. Given the bulk of theory above, it is a legitimate question to ask why this is the case. One possibility is that some defended species suffer too much of a burden from Batesian mimics to sustain a warning signal (see chapter 8). Endler and Mappes (2004) developed a formal model to explore another possibility: that variation among predators might be responsible. Their model assumptions were based on two observations. First, prey are not always predated on by the same species of predator. Second, not all prey defences are effective against all predators. Their model showed that when the frequency of predators that were immune to the defence of a prey species was set above a certain threshold, then defended prey only evolved mildly conspicuous colourations. Thus, variation in predator resistance to prey defences can affect the evolution of warning signals. Perhaps a useful extension of their model would allow the evolution of prey defences, and subsequent predator coevolution, to look at trade-offs between (multiple) chemical defences and their relative costs and benefits.

Recently, the role of multi-modality in the evolution of warning signals has been brought to attention. As (Joron, 2003) points out, “Assuming that predators would rely soley on color and not behavior, motion, or sounds is perhaps simplistic, and it is sometimes argued that multiple signals could even be a prerequisite for the evolution of warning coloration.” The reasons behind the importance of multi-modal signals can likely be found in studying the coevolution of predator biases and warning signals. Joron (2003) suggests that multi-modality may have evolved in defended species in an attempt to escape Batesian mimicry from undefended species (see also Chapter 8). Thus, a model of the coevolution of predator biases and multi-modality in prey would be revealing.

4.7 Summary

This chapter has introduced much of the current theory behind the evolution of natural warning signals. Although warning signals have interested scientists since Darwin’s time, there is still much scope for development of theories and there is still much controversy regarding the origin and maintenance of warning signals. Although two alternative perspectives on the subject were discussed they are not necessarily mutually exclusive.

There are many factors influencing the evolution of warning signals, and our best hope in understanding them is to develop simple models and incrementally add more factors. Ruxton et al. (2004a) envision the way ahead: “We suggest that coevolutionary questions will be a major focus of work in aposematism over the next decade.” Evolutionary simulation models, in concert with mathematical models, seem well placed for such a task.

Chapter 5

Warning Signals and Predator-Prey Coevolution

Theories of the evolution of warning signals are typically expressed using analytic and computational models, most of which attribute aspects of predator psychology as the key factors facilitating the evolution of warning signals. Sherratt (2002a) provides a novel and promising perspective with a model that considers the coevolution of predator and prey populations, showing how predators may develop a bias towards attacking cryptic prey in preference to conspicuous prey. Here an individual-based replication of the model is presented, and the results find, in accordance with Sherratt, that predators evolve a bias towards attacking cryptic prey. A Monte Carlo simulation is then used to calculate the relative survivorships of cryptic and conspicuous prey and stress that, as it stands, the model does not predict the evolution or stability of warning signals. The model is extended by giving predators continuous attack strategies and by allowing the evolution of prey conspicuousness: results are robust to the first modification but, in all cases, cryptic prey always enjoy a higher survivorship than conspicuous prey. When conspicuousness is allowed to evolve, prey quickly evolve towards crypsis, even when runaway coevolution is enabled. Sherratt's approach is promising, but other aspects of predator psychology, besides their innate response, remain vital to our understanding of warning signals.

5.1 Introduction

As noted in Chapter 1, nature is rich with organisms that display bright warning colourations. Although such displays intuitively appear to be an honest indication of a defense mechanism, biologists have long puzzled over their evolutionary origins. Specifically, most studies concentrate on finding conditions under which defended organisms will evolve conspicuous, rather than cryptic, colorations. Why, for example, do bees flaunt bright stripes when a more cryptic form would help to hide them from predators? In particular, if we assume that ancestral bees were cryptic, what was the evolutionary advantage for a conspicuous mutant?

Theories of the evolution of warning signals are typically expressed using analytic and computational models. Most models attribute aspects of predator psychology (e.g., learning and forgetting rates) as key factors facilitating the evolution of warning signals (see, e.g., Harvey et al., 1982; Sillén-Tullberg & Bryant, 1983; Leimar et al., 1986; Guilford, 1990; Yachi & Higashi, 1998; Servedio, 2000; Speed, 2001b). Such models are typically receiver oriented, i.e., they focus on predator response to the presence or absence of warning signals. The models generally only include a single predator, and predict the evolution of warning signals over a relatively narrow range of conditions. For example, Speed (2001b) predicts generally that warning signals will evolve when predators, equipped with the ability to learn, have a better memory for aposematic prey than for cryptic prey, or when predators are neophobic and have some degree of forgetting.

Sherratt (2002a) provides an innovative perspective on the evolution of warning signals by considering coevolving predator and prey populations. The model's predators are deterministic in that they have a fixed behavioral strategy over their lifetime, and cannot learn from experience. For both cryptic and conspicuous prey, each predator has a fixed policy of either attacking or avoiding.

The model stands out from other approaches in several ways: it allows predators and prey to coevolve, it does not primarily rely on predator psychology as an explanatory factor, and it is the first to acknowledge that the decision of one predator can, through evolution, influence the decisions of future predators. Thus, it marks a shift in emphasis for warning signal research to examining predator-prey coevolutionary approaches. However, the ambitious aims of Sherratt's work coupled with its mathematical framework have necessarily lead to a focus on the selective pressures affecting predators rather than prey. This leaves room for an individual-based model in which both sides of the coevolutionary relationship can be thoroughly addressed.

Here, a description of Sherratt's model is first presented. Second, an individual-based simulation model is presented as a reworking of his analytic model, and we consider prey survivorship over a range of conditions in order to further explore the situations under which conspicuous colorations are evolutionarily advantageous. Extensions to the model are then discussed in which:

- Predators have continuous attack strategies rather than being restricted to two or three behavioural options.
- Prey conspicuousness can evolve.
- 'Runaway coevolution' is enabled, by dropping the assumption that migrating prey will be drawn from a fixed distribution of cryptic versus conspicuous and undefended versus defended individuals.

5.2 Sherratt's Model

Sherratt's (2002a) model assumes a world where a diverse range of prey migrate to a locality inhabited by a population of predators. Prey are assigned a level of conspicuousness measured by p—their probability of detection by predators. In the analytic model, all possible values of p are considered, whereas in a subsequent simulation version of the model, prey are either highly conspicuous ($p = 0.9$) or highly cryptic ($p = 0.1$). This detectability parameter represents how easily a prey item can be perceived as distinct from its background (Guilford, 1990). An individual prey item entering the locality is defended with probability $1 - q$ or undefended with probability q (usually $q = 0.8$). The key variables considered in both versions of the model are the probabilities of prey of a particular level of conspicuousness being either defended or undefended. These probabilities are influenced by the effects of predation over time, and thus we are shown, for example, that conspicuous prey are more likely to be defended than are cryptic prey. Effectively, prey defense levels can evolve, but prey conspicuousness (and thus warning signals *per se*) cannot.

Defended and undefended prey survive predatory attacks with probabilities s_d and s_u respectively (where $s_d > s_u$). Prey are assumed to live long enough to breed if they survive predators searching 'in their vicinity' on t occasions. Predators are offered prey at random and, assuming that the prey is detected, will choose whether or not to attack depending on their innate strategy for prey of conspicuousness p. In the basic model, x_p predators attack prey of conspicuousness p and y_p predators do not ($x_p + y_p = n$). Predators are rewarded with a fitness benefit b for attacking an undefended prey item, and receive a fitness deduction c for attacking a defended prey item (where $c > b$).

Sherratt built an analytic model around the above assumptions and found the evolutionarily stable strategy (ESS) for predators. Under a range of conditions, the model predicts that all predators should attack cryptic prey on encounter, whereas a mixture of attacking and non-attacking strategies is typically predicted for conspicuous prey. These results are due to the following factors:

1. In the model, defended prey are more likely to escape predators, and thus are more likely to survive and reproduce in both cryptic and conspicuous populations.
2. Predators detect more conspicuous prey than cryptic prey, and thus filter out more undefended prey from conspicuous populations. In combination with (1) above, this results in a correlation between prey conspicuousness and defence levels; it is unsurprising that future generations of predators will exploit this information.

In a refinement of the model, an additional strategy was enabled where z_p predators could cautiously attack prey ($z_p + y_p + x_p = n$). This strategy, although resulting in a higher escape probability for both defended and undefended prey (sc_d and sc_u respectively), brought a lower cost $h \times c$ to predators for cautiously attacking defended prey. This model was implemented as a computer simulation in which predators were represented in an individual-based manner. As in the basic model, this model predicts that all predators should attack cryptic prey. However, the refined model predicts that, when encountering conspicuous prey, some predators will not attack at all whereas most will attack cautiously. The resulting predation rates in this refined model are similar, then, to those of the basic model in which roughly half of the predatory population would attack conspicuous prey. Results in both versions of the model can be interpreted as reflecting the tendency of real predators to 'go slow' on conspicuous prey items (Guilford, 1994).

Sherratt uses the model to convey an important point: that the psychological biases of predators may not just be the result of secondary effects of predator nervous systems, but may themselves have been moulded by natural selection. To quote Sherratt (2002a, p. 745) "…I question whether defended prey have tended to evolve conspicuousness simply because that happens to have been the type of signal that predators *a priori* find easiest to learn."

5.3 Method

Individual-based replication of the model described above is now presented. The previous model is extended with a Monte Carlo simulation which calculates the relative survivorships of cryptic and conspicuous prey over a range of parameters; these results may have implications for theories of the two phases of the evolution of warning signals.

5.3.1 Evolutionary Simulation Model

The analytic model was reworked as an individual-based evolutionary simulation, in which predators and prey are represented discretely. The model was built around the same assumptions as Sherratt's model with minor modifications as follows:

- Predators were presented with prey stochastically, with prey items selected using a roulette-wheel selection algorithm acting on prey conspicuousness (i.e., a prey item of conspicuousness 0.9 was nine times more likely to be spotted by a predator than was a prey item of conspicuousness 0.1). This effectively combines the two assumptions of random prey dispersal and higher detectability for conspicuous prey.

- The order in which predators were presented with prey was randomized after each prey generation in order to avoid any artefacts that might emerge from a deterministic selection algorithm.

- Prey randomly reproduced if they were still alive after all n predators finished t foraging attempts (as opposed to surviving if a predator searched in their locality on t occasions).

- Predator generations lasted an order of magnitude longer than prey generations (generation lengths needed to be explicitly encoded in the model and this seemed a reasonable assumption based on real predator-prey systems).

- The parameter t, representing the number of times a predator encounters prey per *prey* generation, was increased from 10 in Sherratt's model to 80 in our own (note that t is defined differently in our model). Such that a predator will experience 800 predation opportunities during its lifetime.

In line with Sherratt's focus on comparisons between defended and undefended prey within conspicuousness classes, prey population sizes were kept constant at 2000 cryptic prey and 2000 conspicuous prey after each generation. Predator strategies for dealing with conspicuous and cryptic prey were represented as two binary loci, with a one encoding attack and a zero encoding avoidance in each case. Costs and benefits for attacking the different types of prey needed to be explicitly encoded; parameter settings used were mainly $c = 2$ and $b = 1$ as suggested by Sherratt in an example (other hard-coded parameters, excepting the value t mentioned above, were also set to Sherratt's example values: see Fig. 3 in his paper). The migration rate m, in our model, represents the actual number of migrating prey per generation. Predator fitness was simply the sum of the costs and benefits experienced by an individual over its lifetime. Successful predator strategies were selected for reproduction using tournament selection (allowing the higher-fitness member of a randomly selected pair of predators to reproduce), and offspring strategies were randomly mutated with a probability of 0.01 per loci.

Each simulation run lasted 1000 prey generations. For each run the mean frequency of each possible predator strategy over the last 200 prey generations was recorded. For each set of parameter values, mean statistics were calculated over 50 runs with different random seed values. In addition to a set of primary runs in which Sherratt's analytic model was duplicated as closely as possible, further runs were carried out in order to systematically vary key parameters, with particular attention to the rate of migration m. The effects of varying other parameters such as t, c, and b were also examined.

5.3.2 Probability of Prey Survival

Before investigating a genuinely coevolutionary model in which prey survivorship is captured explicitly, we wanted to look at how consideration of the question of relative survivorship would affect the results from our individual-based replication of Sherratt's model. Sherratt's model and our replication both give predicted distributions of prey defence levels within a conspicuousness class, and of predation strategies for classes of prey. Given these distributions, a solution to the survivorship question could be found analytically using a hyper-geometric distribution without replacement. However, for simplicity we chose to build a Monte Carlo model. The model took prey defence level distributions and predator strategy distributions from the output of our previous simulation and calculated the mean expected survivorship for conspicuous and cryptic prey over half a million runs. Prey survival was quantified by calculating the percentage of the total population of prey with conspicuousness p surviving each prey season (as in Speed, 2001b). The results were then plotted to determine whether or not the tendency of predators to always attack cryptic prey and to sometimes attack conspicuous prey actually results in a higher survival probability for conspicuous prey.

5.3.3 Evolving Prey Conspicuousness

Although prey defense levels could be said to evolve in the analytic model, and in our individual-based replication of it, the levels of conspicuousness in the prey population were fixed. To further examine whether we should expect the evolution of warning signals under the set of assumptions explored by Sherratt, we extended our simulation model so that one of the two prey populations could explicitly coevolve its level of conspicuousness: each individual in the simulated prey population carried a gene representing its own value for p. This model allows us to directly address the question of whether a prey population in the model could ever be expected to evolve conspicuous coloration despite the obvious benefits of crypsis. In the simulation one population remains cryptic and the other population starts with a conspicuous coloration; we look at whether the latter population will remain conspicuous over evolutionary time.

In this extended model, mutation could act upon the conspicuousness level p of a newborn prey individual with probability 0.01. The effect of mutation was to randomly add or subtract 0.1 from the parental value of p. Minimum and maximum values for p were set at 0.1 and 0.9 respectively. Edge effects in mutation were handled using the reflect operator (Bullock, 1999): if $p < 0.1$ then $p = 0.2$ and if $p > 0.9$ then $p = 0.8$. In Sherratt's original model, prey defence levels change due to the migration of a new prey type. We kept this feature, although it can be regarded as equivalent to mutation of a gene for a prey's defence level. In our model, migration was therefore implemented by periodically killing a random prey item and replacing it with a new prey item of the same conspicuousness, but with a randomly selected defence level. Predators were equipped with strategies specifying whether or not they would attack prey of each of the nine possible levels of conspicuousness. This 'strategy table' for a predator was inherited and the binary entries in the table could be mutated with a probability of 0.01 per locus. All prey in the evolving population were initially highly conspicuous with $p = 0.9$, and the distribution of prey defence levels was taken from randomly selected final generations in the previous simulation. The evolutionary aspect of the model is used to look at whether warning coloration will be stable over time. Initial strategies for predators were randomly determined. The final mean level of conspicuousness for the prey was recorded over the last 200 generations (of 5000 total generations) across 50 runs.

One way in which Sherratt suggests his model can successfully predict the evolution of warning signals is by enabling 'runaway coevolution'. Prey migrating into the modelled environment must, of course, come from somewhere else. However, in these alternative locations, they are likely to be under the same kind of selection pressures. Thus, over time the proportion of undefended conspicuous prey amongst the total migrant intake will decrease as this unfortunate prey type is filtered out in all possible locations. We have therefore implemented an abstraction of this process in the model. Runaway coevolution was implemented by allowing a separate probability for conspicuous and cryptic populations that a migrating prey item is undefended (q_p). If we are to assume that similar predation is occuring outside of the modelled locality, then we can implement this by setting q_p equal to the proportion of undefended prey in the current population each generation. Sherratt suggests that the inclusion of the cautious attack strategy for predators will foster the evolution of warning coloration in general; we therefore enabled this strategy option in our implementation of runaway coevolution.

5.4 Results

We begin by considering the results of our individual-based replication of Sherratt's analytic model. Figure 5.1 displays a typical run, and shows the frequency over generational time of predators with attack strategies for cryptic (a) and conspicuous (b) prey in the basic version of the model in which only attack and no-attack strategies were enabled. Predators evolved to attack cryptic prey but, in the case of conspicuous prey, the population cycled between attack and no-attack strategies.

Figure 5.2 shows the frequency over time of each predatory strategy for cryptic (a) and conspicuous (b) prey over a typical run, where predators could also evolve to attack cautiously. Predators evolved mainly attack strategies for cryptic prey and cycled between no-attack and cautious-attack for conspicuous prey. Figure 5.2 shows results that, at first glance, appear somewhat different to the results in Sherratt's Figure 3; however, the basic story is the same. In both cases, predators evolve to a near-universal attack strategy for cryptic prey, and are more cautious about conspicuous prey.

Migration rate was a key parameter in Sherratt's model. Figure 5.3 shows the mean number of predators with attack strategies for cryptic and conspicuous prey, plotted over a systematic variation of the migration rate. In this case, predators were allowed only attack and no-attack strategies—we find that, over a wide range of migration rates, more predators evolve an attack strategy for cryptic prey than for conspicuous prey. This result shows that Sherratt's finding (that predators are more likely to attack cryptic than conspicuous prey) is robust to variation in the migration rate. When the cautious-attack strategy is enabled (Figure 5.4), predators are still more cautious with conspicuous prey than with cryptic prey under a wide range of migration rates. Similarly, this result is in line with Sherratt's conclusions. Additional sensitivity analyses, not reported here, showed that Sherratt's basic findings were also robust to significant variation in parameters such as the

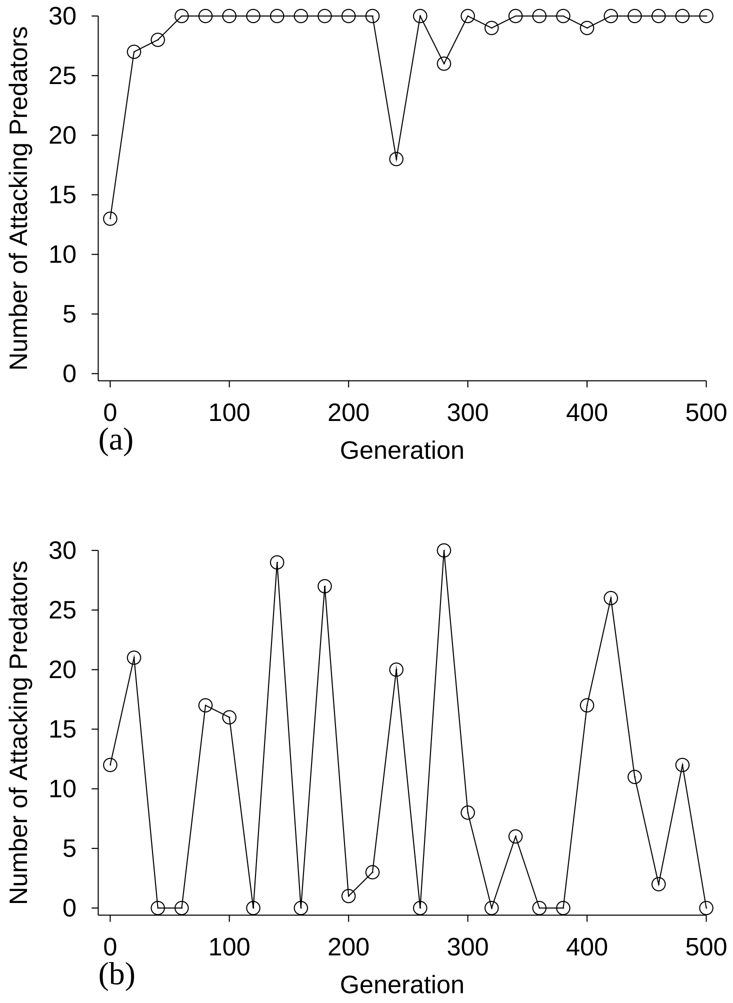

Figure 5.1: Plot of the number of predators with the attack strategy for cryptic prey (a) and conspicuous prey (b) over the first 500 generations of a typical run, where predators were allowed only the no-attack and attack strategies (every 20th generation plotted). Predators evolve attack strategies for cryptic ($p = 0.1$) prey and cycle between no-attack and attack for conspicuous ($p = 0.9$) prey. In this case $n = 30$, $t = 80$, $b = 1$, $c = 2$, $q = 0.8$, $s_u = 0.1$, $s_d = 0.2$, $m = 400$, and $q = 0.8$.

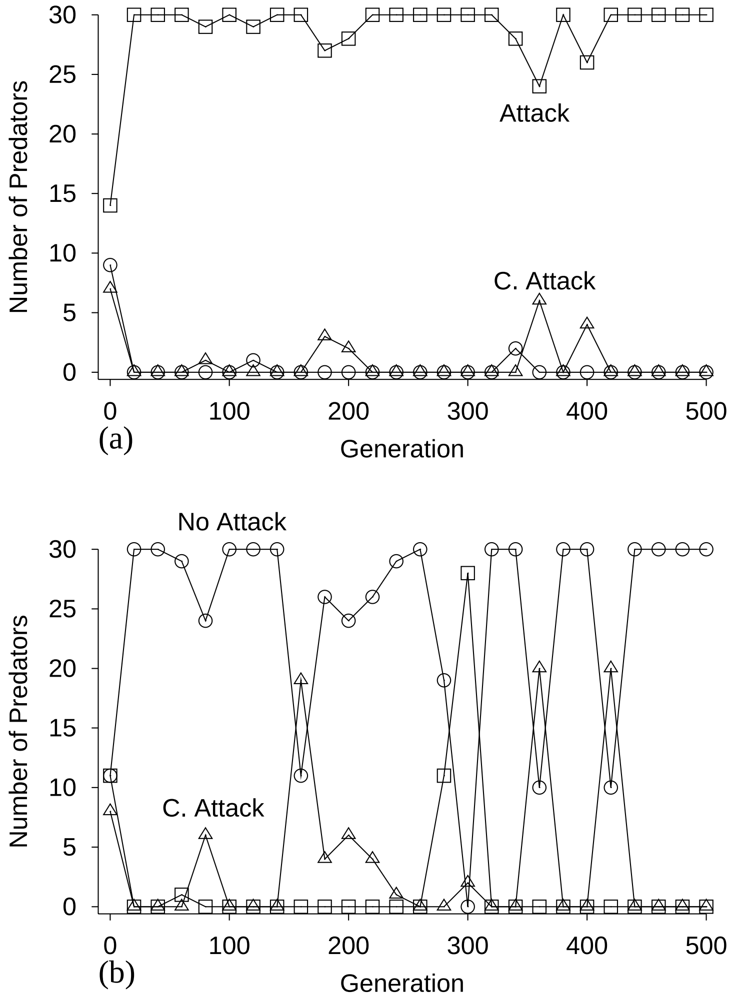

Figure 5.2: Plot of the number of predators with each strategy for cryptic prey (a) and conspicuous prey (b) over the first 500 generations of a typical run, where predators were allowed the cautious attack strategy in addition to the no-attack and attack strategies. Predators evolved attack strategies (squares) for cryptic ($p = 0.1$) prey and typically evolved cautious-attack strategies (triangles) and no-attack strategies (circles) for conspicuous ($p = 0.9$) prey (every 20th generation plotted). In this case $n = 30$, $t = 80$, $b = 1$, $c = 2$, $q = 0.8$, $s_u = 0.1$, $s_d = 0.2$, $sc_u = 0.4$, $sc_d = 0.9$, $h = 0.5$, $m = 400$, and $q = 0.8$

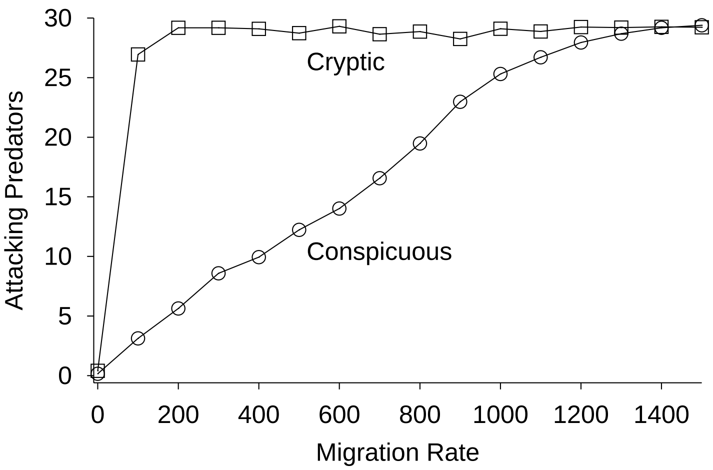

Figure 5.3: The average number of predators with an attack strategy for highly cryptic ($p = 0.1$) prey (squares) and highly conspicuous ($p = 0.9$) prey (circles), where predators were allowed only the no-attack and attack strategies. Results are plotted against various migration rates and are averaged over the last 200 prey generations of 50 runs. In this case $n = 30$, $t = 80$, $b = 1$, $c = 2$, $s_u = 0.1$, $s_d = 0.2$, and $q = 0.8$.

number of predator sampling events, the costs and benefits of predation, and so forth.

We took the analysis of the simulation further, using a Monte Carlo model to calculate the average differential survivorship across conspicuous and cryptic prey. The results of the Monte Carlo model are shown in Figure 5.5; we find that under a wide range of migration rates *cryptic* prey are more likely to survive than conspicuous prey. Monte Carlo runs looking at the difference between defended conspicuous and defended cryptic, and undefended conspicuous and undefended cryptic, showed almost identical results: that both defended and undefended prey had a higher survival probability when cryptic. Supporting simulation runs were conducted in which we varied other parameters such as the costs and benefits of predation, and the escape probabilities for defended and undefended prey (graphs not shown). We found that the model consistently resulted in cryptic prey items enjoying higher survival probabilities than conspicuous prey items. Thus, the advantage to conspicuous prey of predators being less likely to attack them is outweighed by the disadvantage of being easily detected.

When we extend the model by giving predators a continuous value representing their attack probability, we find consistent results. Figure 5.6 shows a typical run where predators evolve a high probability of attacking cryptic prey and a general ambivelance towards conspicuous prey.

Finally, we turn to the results for our explicitly coevolutionary models. When prey conspicuousness p is allowed to evolve we find that, under the conditions tested, the prey

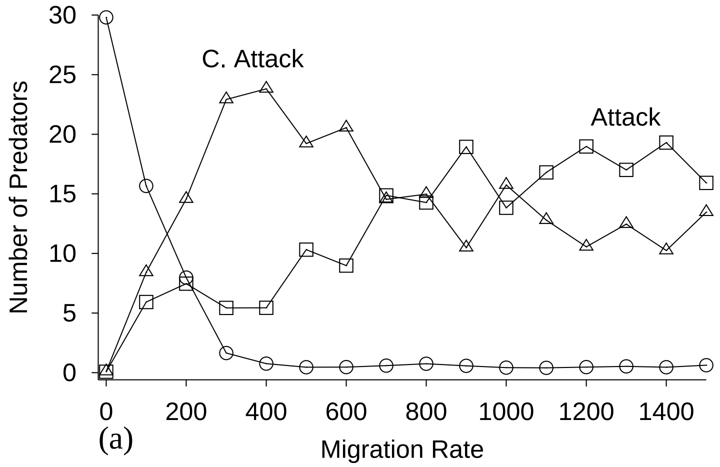

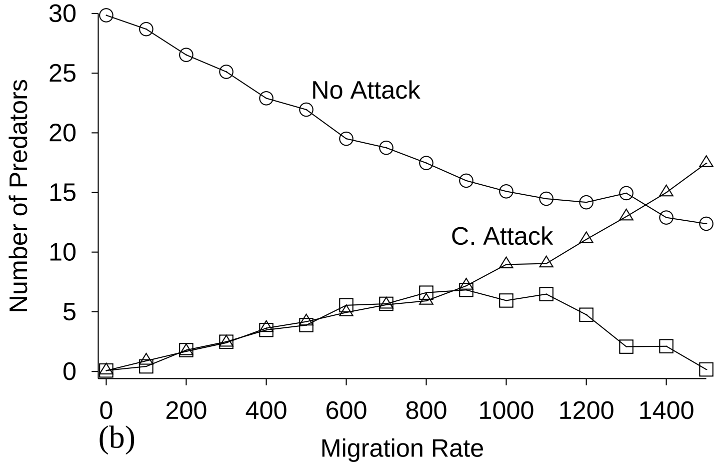

Figure 5.4: The average number of predators with each type of strategy for highly cryptic ($p = 0.1$) prey (a) and highly conspicuous ($p = 0.9$) prey (b), where predators were allowed the cautious attack strategy in addition to the no-attack and attack strategies. Results are plotted against various migration rates and are averaged over the last 200 prey generations of 50 runs. In this case $n = 30$, $t = 80$, $b = 1$, $c = 2$, $s_u = 0.1$, $s_d = 0.2$, $sc_u = 0.4$, $sc_d = 0.9$, $h = 0.5$, and $q = 0.8$.

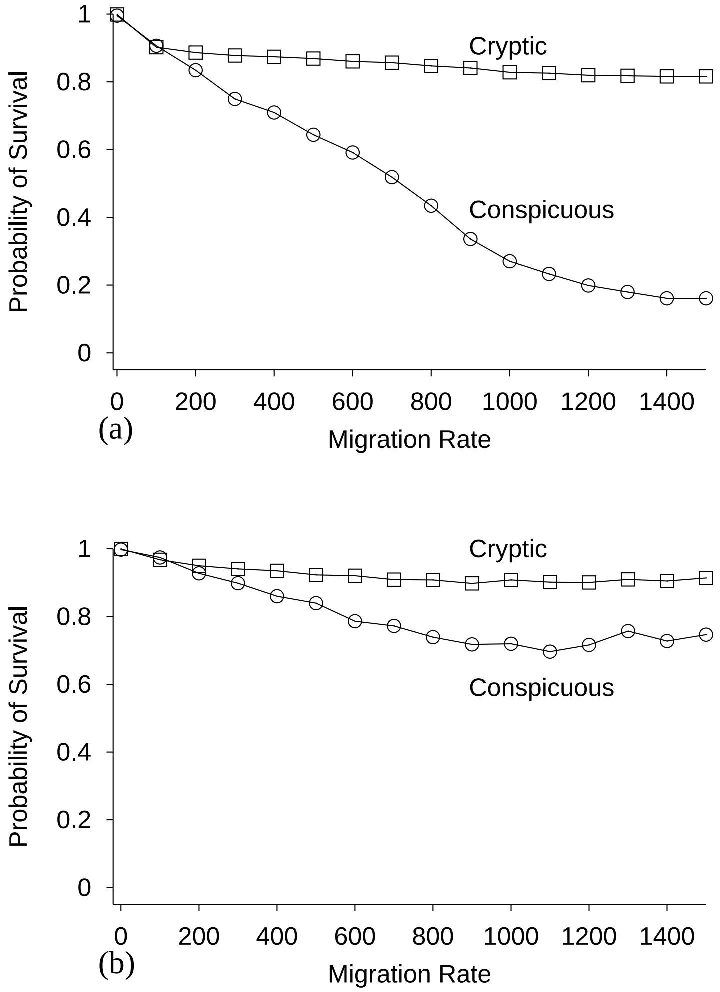

Figure 5.5: The survival probabilities of cryptic ($p = 0.1$) prey and conspicuous ($p = 0.9$) prey, plotted against various migration rates. Both when predators had two (a), or three (b), possible strategies then cryptic prey were typically more likely to survive than conspicuous prey. Results were averaged over half-a-million runs. In this case $n = 30$, $t = 80$, $b = 1$, $c = 2$, $s_u = 0.1$, $s_d = 0.2$, $sc_u = 0.4$, $sc_d = 0.9$, $h = 0.5$, and $q = 0.8$.

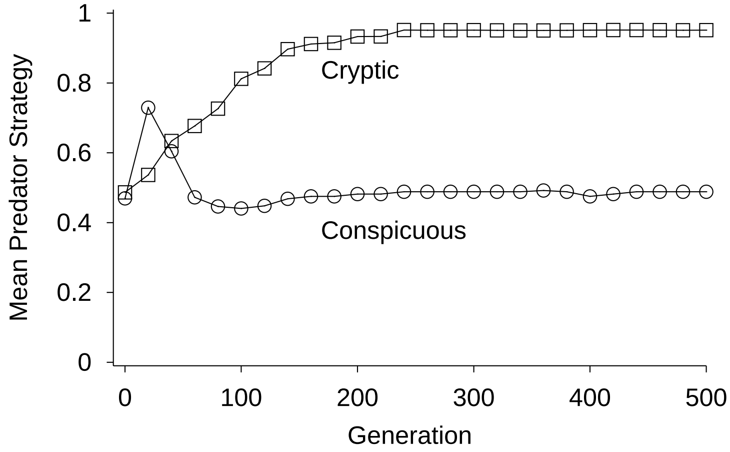

Figure 5.6: The mean strategy for highly cryptic ($p = 0.1$) prey (a) and highly conspicuous ($p = 0.9$) prey (b) in a typical run, where predator strategies were represented as an attack probability in the range [0,1]. In this case $n = 30$, $t = 80$, $b = 1$, $c = 2$, $s_u = 0.1$, $s_d = 0.2$, $h = 0.5$, and $q = 0.8$.

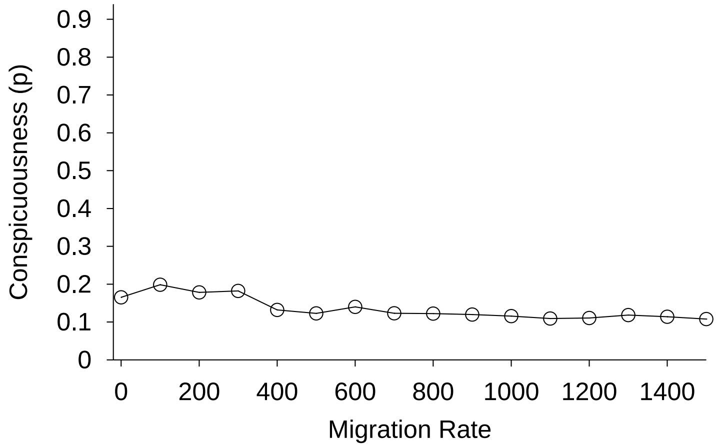

Figure 5.7: The evolved average conspicuousness of prey plotted against various migration rates. Whether predators had two or, in this example, three possible strategies prey that were initially conspicuous evolve to become cryptic. Results were averaged over 20 runs. In this case $n = 30$, $t = 80$, $b = 1$, $c = 2$, $s_u = 0.1$, $s_d = 0.2$, $sc_u = 0.4$, $sc_d = 0.9$, $h = 0.5$, and $q = 0.8$.

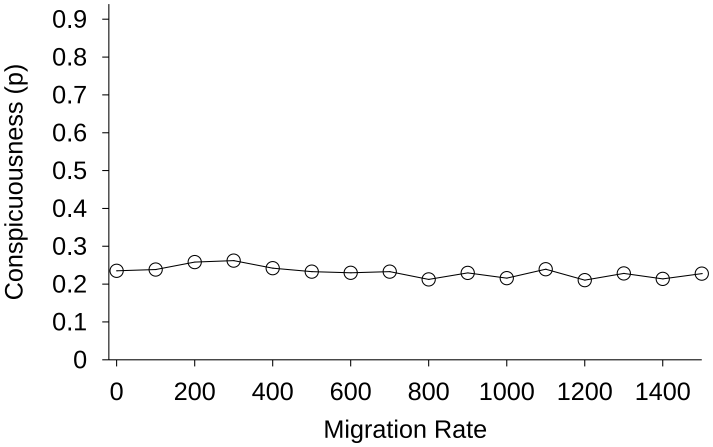

Figure 5.8: The evolved average conspicuousness of prey plotted against various migration rates. When predators were given three possible strategies and runaway coevolution was enabled, prey that were initially conspicuous evolve to become cryptic. Results were averaged over 20 runs. In this case $n = 30$, $t = 80$, $b = 1$, $c = 2$, $s_u = 0.1$, $s_d = 0.2$, $sc_u = 0.4$, $sc_d = 0.9$, $h = 0.5$, and initially $q = 0.8$.

population always evolves to become highly cryptic, despite the initial population being highly conspicuous (Figure 5.7). The results from the Monte Carlo model strongly suggest that evolution will favor crypsis; the results from the coevolutionary model confirm it.

When runaway coevolution is enabled (in addition to enabling the cautious attack strategy), we still find that the conspicuous population evolves towards crypsis. Interestingly, the evolving population converges on $p = 0.2$, rather than the minimum conspicuousness level of $p = 0.1$, in order to distinguish itself from its sister population of fixed cryptic prey, which would contain a much higher frequency of undefended individuals (Figure 5.8).

5.5 Discussion

Our individual-based replication of Sherratt's analytic model confirms that under a wide range of conditions, a general predatory bias to attack cryptic prey more readily than conspicuous prey is predicted. The explanation for this result involves several steps. First, predators filter out more conspicuous prey than cryptic prey, due to the former type's higher probability of detection. Next, as defended prey have a higher chance of escape than undefended prey, predators act to filter out more undefended prey from the conspicuous population than from the cryptic population. This results in a correlation between conspicuous coloration and high defence levels. Finally, the correlation is exploited by the evolving predators and results in their preference for attacking cryptic prey.

The results of the individual runs shown (Figures 5.1 and 5.2) generally match those of Sherratt's analytic model. One minor difference is that, whereas the analytic model predicted that roughly half of the predatory population would attack, our model predicts a cycle between all predators attacking and all predators not attacking. Sherratt makes the common assumption that an equilibrium exists in his models (even in his simulation the best predator is copied over the wost). We do not make this assumption and use a tournament selection algorithm. Thus, the cycles are likely to be due to evolved responses to fluctuations in frequencies of undefended prey; predators learn to attack the prey and consequently filter prey until mostly defended prey are left, then predators learn to avoid the prey and no longer filter the migrating undefended prey which spread within the population. This process repeats, resulting in the cycles. Despite the cycles, the models give the same predictions on average.

We further scrutinized the simulation by constructing a Monte Carlo model of prey survivorship. The results from this model predict that, under the assumptions and conditions stipulated by Sherratt, cryptic prey will enjoy higher survival rates, and the evolution of warning signals should *not* be expected. Sherratt recognised this limitation for his analytical model, pointing out that the evolved behavioural responses of predators to novel prey cannot in itself explain aposematism, but may help explain why it is not so strongly selected against. This conclusion is further supported by an extended simulation model, in which prey conspicuousness could explicitly evolve. Under the conditions tested, prey always rapidly evolved to be cryptic. Note that we look at relative survivorship across cryptic and conspicuous prey and even though the cryptic population contains a higher frequency of undefended prey than the conspicuous population, we still show that cryptic prey enjoy higher survivorship. This means that conspicuous prey suffer a lower survivorship *despite* having an extra survival advantage in their likelihood of escaping from predators. Thus, this consideration raises an even bigger question mark over why defended prey would ever do anything other than evolve towards crypsis.

Even though our replication of Sherratt's analytic model confirms his predictions, unfortunately the additional analyses we have conducted suggest that his results do not lead to the desired conclusions in an evolutionary sense. The observed predator bias towards attacking cryptic prey does not result in a selective advantage for conspicuous prey once differential survivorship is fully taken into account. This would suggest that theories invoking predator psychology, e.g., the need for naïve predators to avoid costly mistakes when selecting prey, are still currently the most likely candidates for workable explanations for the fixation of warning signals in a population (Speed, 2001b, offers a good example of a relevant model). However, it is probable that future models incorporating predator psychology *and* predator-prey coevolution will lead to a more complete picture of the evolution of warning signals (but see Chapter 8 for an alternative approach).

Why might the model presented in (Sherratt, 2002a) cause confusion as to its implications? Sherratt's model allows for two fixed populations: one conspicuous and one cryptic, each of which can evolve varying levels of defence. However, the typical approach is to model two populations where one is defended and the other undefended. Thus, the evolving trait is their level of conspicuousness/crypsis. Thus Sherratt's model can almost be looked at as the evolution of defence in already conspicuous prey, and the evolution towards no defence in cryptic prey. As Sherratt's model always results in conspicuous defended prey and undefended cryptic prey (and predators that disproportionally attack cryptic prey), it is easily misinterpreted as explaining the evolution of warning signals. If the level of conspicuousness is allowed to evolve along with the level of defence then, in the model, both defended and undefended prey evolve to become cryptic. Thus, the work here has made clear the implications of the model.

Sherratt (2002a) suggested that, within the framework of his model, one of the most likely scenarios for the evolution of warning signals would be runaway coevolution and the presence of the cautious attack strategy. However, when we explicitly modelled the combination of these two phenomena, we still found that the evolving population tended towards crypsis. The fact that the evolving population clustered around $p = 0.2$ shows that the decision-making process of predators, represented in our model by a simple nine-element strategy, had an effect on the evolution of prey: the evolving population was caught between a drive for crypsis and a need to remain distinct from the fixed population ($p = 0.1$) which suffered a higher predation rate. Sherratt (2002, p. 745) himself notes that "…learning by direct experience probably plays an extremely important role in facilitating the evolution of aposematism." Given our results it certainly appears that the psychological properties of predators (generalization, learning, forgetting, and so forth), or other costs to crypsis, may prove vital to our understanding of the evolution and maintenance of warning signals. However, Sherratt's approach, and our additional simulations, suggest that the biases observed in receiver psychology might not be fixed, as is commonly assumed, but may be subject to selection.

Speed and Ruxton (2002) discuss Sherratt's model and suggest that further simulation work is needed to explore the evolution of warning signals under different conditions. Instead of assuming a world that is inhabited by a diversity of defended and undefended prey types, what if we drop the assumption of large amounts of diverse prey migration and assume a world where conspicuous mutants arise only rarely? Our supporting simulation work has shed light on this question. When migration is dropped from the simulation, and occasional prey defense mutation is implemented (in addition to conspicuousness mutating) as with the genuinely coevolutionary model, we find that prey evolve to become cryptic (as in Figure 5.7). Thus, we find that when the assumption of a world inhabited by a highly diverse range of prey types is dropped, prey should also be expected to evolve to become cryptic under the assumptions outlined in the model.

The significant contribution of Sherratt's work is to move theoretical studies of the evolution of warning signals towards considering coevolving populations of predators and prey. This could prove to be a significant advance, as most theoretical studies of warning signals and mimicry consider a lone predator. Though this approach does not explain the survival to fixation of warning signals, it introduces an innovative and promising route to exploring aposematic phenomena.

Chapter 6

Mimicry: A classical coevolutionary paradigm

6.1 Introduction

The adaptive resemblance of one species to another is termed *mimicry*. This is a textbook example of evolutionary adaptation as it is initially intuitive and appealing. As such, it is discussed as a key example in most evolutionary biology textbooks. Mimicry studies date to a few years after the publication of *On the Origin of Species*, and mimicry was presented as compelling evidence for evolutionary theory (Wallace, 1866). However, Turner (1995, p. 131) states that "…although mimicry is probably one of the most thoroughly explored exemplars [of coevolution] it is surprising what we do not know about it." Mimicry is an interesting case of coevolution, as the model (the species being mimicked) and mimic do not interact directly; instead they are tied to one another through the actions of their predators. Ronald Aylmer Fisher, who almost single-handedly created the foundations for modern statistical science, was interested in warning signals and mimicry. Fisher (1930) provided early discussions of warning signals and mimicry, and used them as evidence for the theory of evolution in general. He called the theory of mimicry "…the greatest post-Darwinian application of Natural Selection."

Many different types of mimicry exist in nature. Examples include aggressive mimicry (a wolf in sheep's clothing affair), pollinator mimicry (where, for example, flowers attract bees by mimicking the appearance of a female bee), and sexual mimicry (where males mimic the appearance of females in order to get close to them), among others. These examples of mimicry are of inherent interest; however, this thesis is concerned solely with interactions between predators and prey and the classic forms of mimicry such as the Batesian and Müllerian varieties.

6.2 Batesian Mimicry: Parasitic mimicry

Undefended prey are vulnerable to predation. As we have seen, one way that their survival chances have been enhanced is with the evolution of crypsis. However, there is another way that undefended prey can gain protection against predators—by mimicking the appearance of defended prey. A species displaying warning signals is sending an honest signal to predators that it is defended. However, the colouration typically used for signalling a warning can evolve regardless of the quality, i.e., the level of defence, that is being signalled. Thus, the honest signal given by defended prey leaves itself open to bluffers; some undefended prey evolve conspicuous colourations such that predators struggle to distinguish them from defended prey. It may not be enough to just evolve a conspicuous colouration, as predators would still learn to distinguish the undefended prey from the defended prey eventually (but still with some errors in recognition). Thus, undefended prey need to evolve as close a resemblance to defended prey as possible (i.e., in pattern, colour, and so forth, as well as conspicuousness), so that they cannot be distinguished by the predator in question.

The hoverfly *Chrysotoxum festivum*, for an example, exhibits a black and yellow striped warning colouration despite being palatable to birds, and lacking a sting. This type of mimicry can be described as a parasitic relationship (see, e.g., Wickler, 1968) as the *model* species, in this case bumblebees or wasps, suffers from a dilution of the aversive effects of its colouration (Plowright & Owen, 1980). The defended species would suffer because some of the prey items sharing its appearance would provide a positive experience for attacking predators—increasing the chance that predators would mistakenly attack the defended species. The mimic, however, gains from the bluff and enjoys reduced predation without the need to evolve a costly defence such as a sting or toxin. This is known as Batesian mimicry after the discoverer of its adaptive value: Bates (1862). Batesian mimicry can be defined as *a parasitic mimetic relationship between a palatable and a defended species*. Here, parasitic refers to a situation where one species enjoys increased fitness at the expense of the fitness of another species. Batesian mimicry is bad for predators as potential meals are presented as defended. Empirical evidence for

Batesian mimicry is extensive (Nicholson, 1927; Brower, 1958a, 1958b, 1958c, 1960); see Ruxton et al. (2004a) for a review.

6.3 Müllerian mimicry: Mutualistic mimicry

Many different species of *Heliconius* butterflies share a conspicuous colouration, making an honest display as to their anti-predator defences (Turner, 1981; Sheppard et al., 1985; Mallet & Gilbert, 1995). This is a case of Müllerian mimicry (Müller, 1879). This mutualistic mimicry is common among (but not between) many species of butterflies and bumblebees, and occurs in vertebrates such as poison arrow frogs (Symula et al., 2001). Müllerian mimicry can be defined as *a mutualistic mimetic relationship between two or more defended species*. Although it is easy to imagine the adaptive advantage afforded to a Batesian mimic, the adaptive advantage of defended species sharing a colouration is less obvious. In this case, defended species that share an appearance each sacrifice fewer individuals to naïve predators during their learning phase. This mutualism is the reason that many defended species evolve a uniform pattern. Predators also typically benefit from a reduced educational burden, as predators treat the two prey species involved as a single species, removing the need to learn to avoid each species separately (Fisher, 1930; Endler, 1991; Ritland, 1991, although see Speed, 1993b). Thus, both prey species in the relationship mutually benefit from reducing the burden of predator education (Fisher, 1930).

6.4 Coevolutionary Dynamics

Because predators generalise, it is possible for them to mistake approximate mimics for their model. Thus mimics typically receive more protection as their mimicry becomes more refined (although mimetic species do also exhibit imperfect mimicry). Predators can discriminate between prey of sufficiently distinct colourations. For mimicry to evolve, then, an initial resemblance (in the eyes of the predator) is needed—a matter that we will return to later.

The coevolutionary dynamics involved in Müllerian and Batesian mimicry differ (Turner, 1987, 1995, see Figure 6.1). Müllerian mimics generally *converge* upon the same colour pattern; selection pressure typically causes both species to evolve ‘towards’ a middle-ground colour pattern.[1] Batesian mimics *adverge* (Brower & Brower, 1972); a conflict of

---

[1] Although this is not *always* true; a new immigrant defended prey species might adverge upon abundant existing defended prey.

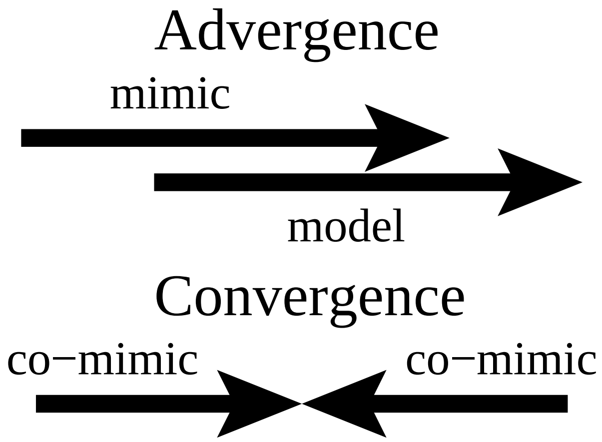

Figure 6.1: The difference in dynamics between Batesian and Müllerian mimicry (adapted from Turner, 1995).

interests between a Batesian mimic and its model results in a coevolutionary arms-race; the mimic ‘chases’ the model through colouration space. Dixey (1909) reasoned that the appearance of the model species should not change as a result of the mimicry, but this was refuted by Fisher (1930) who argued that “selection will tend to modify the model so as to render it different from the mimic.” The Batesian mimic is typically expected to keep up with the model, as the selection pressure on a Batesian mimic to gain protection is generally greater than the selection pressure on a model to evade the mimic. Thus, the model is expected to move away from the mimic, but at a slower rate (Nur, 1970; Turner, 1977, 1987; Sheppard et al., 1985). Brower and Brower (1972, p. 66) state the evolutionary dynamics relevant to this work aptly: “…Batesian mimicry promotes continuous change in time, whereas Müllerian mimicry tends towards stabilization of common color patterns.” For a good model of mimicry dynamics see Gavrilets and Hastings (1998).

6.5 Batesian and Müllerian Mimicry: Contrasts and Comparisons

6.5.1 Similarities and Differences

The difference between Batesian and Müllerian mimicry is often discussed with reference to the differences in levels of defence of the interacting prey species: Batesian mimicry occurs when an undefended species mimics a defended species, and Müllerian mimicry occurs when two defended species mimic each other. Thus, Batesian mimicry can be considered parasitic and Müllerian mimicry considered mutualistic. The result is a difference in dynamics between these forms of mimicry. A further difference comes from the effect of frequency dependence on the different mimetic forms. Whereas Müllerian mimics benefit from high population frequencies, due to there being more defended prey to share the burden of educating the predator, Batesian mimics benefit from rarity (Turner, 1987).

The above points follow if the split between Batesian and Müllerian mimicry is clear. However, the distinction is sometimes blurred. So far prey have been described as either *defended* or *undefended*, when a more detailed (but not always appropriate) account would talk of defence *levels*. One could think of this as a *palatability spectrum*, where one end of the spectrum represents highly defended prey and the other undefended prey (Brower, Ryerson, Coppinger, & Glazier, 1968; Pough, Brower, Meck, & Kessell, 1973; Turner, 1984; Turner, Kearney, & Exton, 1984). To quote Turner (1984) (p. 1) "[The palatability spectrum] is ecojargon for the fact that all things are not equally nice to eat." Although the idea of the palatability spectrum is useful as a simplified conceptual tool, it too is not always fully correct or appropriate. For example, multiple predator species might hunt the same prey species, which might be highly defended from the point of view of one set of predators but only mildly (or not at all) defended from the perspective of the other. This type of effect could obviously alter the dynamics of mimicry as could, for example, differences in predators' cognitive and physical abilities to track or to catch prey (Malcolm, 1990). Ruxton et al. (2004a, p. 169) comment on how this could help to blur the boundaries between Müllerian and Batesian mimicry: "...when predators differ widely in their preferences then one might expect mimicry systems to represent some form of hybrid between classical Batesian mimicry and Müllerian mimicry."

6.5.2 Quasi-Batesian: A Different Mimetic Category?

The argument above could lead us to question whether a particular mimic is actually a Batesian parasite, a Müllerian mutualist, or something in between. Owen and Owen (1984) developed one of the first models to look at this issue in detail. In their mathematical model, predators could learn, and had short memories. [2] Predators were modelled as coming across one of two prey species, which shared a fixed appearance, at random. The levels of defence of each species were examined with reference to predation rates on both species. Owen & Owen found that classical Batesian and Müllerian mimicry can arise in the model under certain conditions. However, if the defence levels of the two species start out equal and become more and more unequal, then the mimetic relationship will shift

---

[2] Owen & Owen's implementation of predator memory is arguably unrealistically short, such that there is a minimum 50% attack rate (Ruxton et al., 2004a) even for a prey type that has been repeatedly found to be defended.

from approximately Müllerian to approximately Batesian in character. Furthermore, the balance between the two types of mimicry will depend on the relative abundances of the two species (for a stochastic replication of this work see Speed, 1999).

More notoriously, Speed (1993a) further questioned the division between Batesian and Müllerian mimicry using a stochastic computer model. Speed (1993a) implemented predation with learning and forgetting algorithms, using ideas from Pavlovian conditioning theory. He showed that when there are differences in levels of defence between two mimetic defended prey species, the lesser defended prey could (depending on conditions) dilute the protection afforded by the more highly defended species. Thus, even a defended mimic, which would conventionally be seen as mutualistic, can be parasitic in some situations. Speed (1993a) termed this *quasi-Batesian* mimicry (see also, Turner & Speed, 1996, 1999b; Speed & Turner, 1999; Speed, 2001a).

A related model (Kokko, Mappes, & Lindstrom, 2003) has shown that the presence of alternative prey could change model-mimic dynamics between mutualism and parasitism, even when the mimic and model share the same level of defence. However, it should be noted that Kokko et al.'s (2003) results are somewhat controversial, as they conflict with the results of a series of models which have examined the effect of introducing alternative (palatable) prey into the system; these models have simply predicted that predator interest will shift to the new species and that there will be a corresponding reduction in the predation rate on both models and mimics (Holling, 1965; Emlen, 1968b; Luedeman, McMorris, & Warner, 1981; Getty, 1985; Sherratt & Beatty, 2003).

MacDougall and Dawkins (1999) interpreted Speed's quasi-Batesian model as predicting that defended prey must share identical defence levels in order to be Müllerian in the conventional sense (i.e., that Batesian or quasi-Batesian mimicry was inevitable given differing defence levels). MacDougall & Dawkins further reasoned that, given the low likelihood that two defended mimetic species would share identical defence levels in nature, one would expect true Müllerian mimicry to be extremely rare. MacDougall & Dawkins disputed this notion, and constructed a modification of Speed's (1993a) model in which predator discrimination errors were now possible. They showed that if predators can make discrimination errors then the classic split between Müllerian and Batesian mimicry holds. MacDougall & Dawkins explain why this occurs: when predator discrimination errors are added, both defended mimetic species benefit from a lower risk of being mistaken for a third, palatable species. Ruxton (1998) considered the significance of this result: "The conceptual leap here is that the mimetic relationship between two species cannot be understood without proper consideration of their shared predator. Whether the relationship is Müllerian or Batesian will depend on the discrimination abilities of the predator, and also on the properties of alternative prey. The relationship could be different for various predators, or for the same predator at different times or in different places." However, Speed and Turner (1999) pointed out that a mistake in MacDougall and Dawkins' understanding of the original model's predictions; the original model predicted Müllerian mimicry not just when prey defence levels are identical, but when they are "…nearly as nasty as their comimics" (Speed & Turner, 1999). Further, a possible problem with the extended model was the exclusion of a forgetting mechanism, which was included in Speed's original model. A future extension of (MacDougall & Dawkins, 1999) is needed where forgetting is included in the model (Speed & Turner, 1999).

The twists and turns of the theoretical debate outlined above indicate that we still have much to learn about the nature of mimetic relationships. The models presented in this thesis will, of necessity, start by examining classical Batesian and Müllerian relationships. However, where possible, we will use simulation models as a tool for looking at the origins of more complex mimetic relationships.

6.6 Models of Batesian and Müllerian Mimicry

6.6.1 The Genesis of Mimicry

How does mimicry get started? This question has been answered by theorists in various ways, and regrettably the current state of our empirical knowledge of mimicry makes it difficult to choose between the competing theoretical accounts. The earliest and most widely cited account of mimicry's origins invokes a two stage process:

- Stage one: after an extensive period of time in which no mimetic relationship exists, a major mutation (such as a mutation on a modifier gene regulating a super-gene complex Turner, 1988) produces a mutant individual that is an approximate mimic by chance (Punnett, 1915; Goldschmidt, 1945).
- Stage two: The first stage is likely to produce poor mimics. However, the familiar machinery of natural selection will then kick in, and the quality of the mimicry will gradually improve (Nicholson, 1927).

It might also be possible for mimicry to get started *without* a dramatic major mutation in a single individual, if genetic drift brought an entire population within some sort of "similarity radius" of a model species, followed by gradual improvement (see chapter 7). Alternatively, two species could be phylogenetically proximate from the outset. Whatever the proposed mechanism for generating an initial degree of resemblance, most models of mimicry tend to concentrate on stage two.

The dynamics of Batesian and Müllerian mimicry differ, but the effect of interactions between the two types of mimicry is under-explored (although see Yamauchi, 1993; Gavrilets & Hastings, 1998). Most current models concentrate on just two prey species and, therefore, one type of mimicry at any one time. Using more complex community models—in which more than two prey species are modelled—we can examine how the different forms of mimicry relate (chapter 7).

6.6.2 Models of Batesian Mimicry

Batesian mimicry has been extensively modelled. However, most models, of both Batesian and Müllerian mimicry, do not allow the mimic, or its model, to evolve. Instead they typically examine the difference in predation rate on defended and undefended prey when either perfect Batesian mimicry is in place, or when the two prey types are distinct (i.e., there is no mimicry). A discussion of a representative model (Oaten, Pearce, & Smyth, 1975) will help to characterise the typical approach and give a flavour of the type of models that are currently used.

There have been a vast number of attempts to model Batesian mimicry. Most of the models have used traditional mathematical approaches to compare predation rates on phenotypes for two species. In such models, there is typically one condition in which the two species are assumed to be exactly alike, and another in which they are completely distinguishable (e.g., Holling, 1965; Estabrook & Jespersen, 1974; Bobisud & Potratz, 1976; Turner et al., 1984; Kannan, 1983). The difference in predation rate across these two conditions is held to be a measure of the benefit of the mimetic relationship to the undefended species. A refinement of this approach has been to incorporate signal detection theory into models of mimicry (Oaten et al., 1975; Getty, 1985; Greenwood, 1986). Signal-detection theory allows investigators to consider varying levels of similarity between models and mimics—in other words, to consider degrees of imperfect mimicry rather than a simple dichotomy between perfect resemblance and complete distinctiveness. In most of these models, predators do not learn, but are assumed to act appropriately to maximise a given set of costs and benefits (see chapter 8).

The predictions drawn from these models of Batesian mimicry have been quite consistent. The rate of predator attacks borne by the model species is expected to increase as the ratio of Batesian mimics to models increases (see also Emlen, 1968a). In addition, most models show (or assume) that evolution should refine a Batesian mimic's colouration until it resembles the model perfectly.

However, there are observations in nature of imperfect Batesian mimics. Several models have offered explanations for this observation, with reasons including genetic constraints (Charlesworth & Charlesworth, 1975), a lack of further selection pressure (Sherratt, 2002b, 2003) and kin selection (Johnstone, 2002). There is another possibility. As we have seen, Batesian mimics and their models are thought to undergo a continuous coevolutionary arms-race in which the model evolves away from the mimic, only for the mimic to evolve towards the model again. Thus, because there is no evolutionarily stable strategy in this case, we may be observing the mimetic relationship during a period of time where the model has evolved away from the mimic; if we came back in the distant future when the mimic has had enough time to catch up, we might see a more perfect mimicry. Whatever the correct explanation, and there may be more than one, it is noteworthy that many mimetic species show imperfect mimicry, and that a growing number of theories (Oaten et al., 1975; Azmeh, Owen, Sorensen, Grewcock, & Gilbert, 1998; Edmunds, 2000; Johnstone, 2002; Sherratt, 2002b) exist to explain this phenomenon.

Oaten et al. (1975) developed one of the simplest models of Batesian mimicry. The model does not attempt to incorporate any predator memory, learning, or forgetting, but uses a signal detection approach to look at the evolution of Batesian mimicry. The model allows different levels of similarity between model and mimic such that the relative costs of different phenotypic strategies can be calculated, and solutions found in terms of a predicted stable distribution of strategies. Oaten et al.'s model specified that predators should attack prey that show a signal (or colouration) s if:

$$bf_1(s)p > cf_2(s)q \tag{6.1}$$

Where b is the benefit and c is the cost of attacking a mimic and model respectively, p and q are the relative frequencies of mimics and models, and $f_1(s)$ and $f_2(s)$ give the probability density of mimics and models with signal s. In addition to giving a simple formula for understanding Batesian mimicry, the model is the first to show conditions for the evolution of imperfect mimicry. The equation helps us understand conditions under which predators might still pay attention to a warning signal, and when the effectiveness of the signal might temporarily break down. (For a similar popular model see Huheey, 1964).

Most theoretical studies of mimicry have focused on strategic factors (rather than, for example, evolutionary dynamics). Analytic and Monte Carlo models have become a common approach to theoretical studies of mimicry. Using a Monte Carlo approach Turner et al. (1984) looked at predation rates on the following different prey types:

- Nasty: A defended prey type with a distinct appearance;
- Solo: An undefended prey type with a distinct appearance;
- Model: A defended prey type sharing an appearance with the mimic;
- Mimic: An undefended prey type sharing an appearance with the model.

Each time-step a single predator was assumed to encounter one of these prey types randomly in proportion to their relative population densities. When the predator encounters a prey type the attack decision is made by checking that a random number is below the attack probability threshold p. After an attack the predator updates an attack probability for the given prey type depending on the prey's defence level. The model and mimic share the same attack probability as the predator is unable to distinguish between the two types. Over time the predator forgets what it has learnt—although this can be countered by subsequent learning. The model was run over many time cycles in order to assess the resulting predator behaviour. The results of the model conform with most theoretical studies of mimicry and can be outlined as follows:

- The fitness of Müllerian mimics is positively density-dependent (they enjoy more protection when common);
- The fitness of Batesian mimics is negatively density-dependent (they enjoy more protection when rare);
- Predators learn to attack Solo (undefended) prey and avoid Nasty (defended) prey;
- Mimics always benefit from their mimicry.

Models such as those of Oaten et al. (1975) and Turner et al. (1984) provide a representative sampling of the efforts that were made in the 1970s and the 1980s to capture the evolution of Batesian mimicry in a formal model. More recent work includes Holmgren and Enquist (1999) who presented a novel approach to models of mimicry which took inspiration from the world of artificial intelligence (as does the approach used in this thesis). The model employed artificial neural networks (multi-layered perceptrons) to represent the response of predators to prey. Both predators and prey could evolve in the model. The prey phenotype was modelled as a nine-dimensional vector the euclidean distance between two phenotypes was used as a measure of (dis)similarity between prey of different species. In line with previous studies, the authors found that Batesian mimics evolve towards the model (they adverge) and in response the model evolves away from the mimic. However, in contrast to previous thinking Holmgren and Enquist show that the model sometimes moves towards the mimic. The authors explain these results as products of network response displacements (p. 155): “Typically, neither the model nor the mimic is located where the receiver is most likely to react with the favoured response. Instead both players are evolving towards the response peak with the model usually being somewhat ahead, i.e., closer to the peak than the mimic.” Thus, if the network’s response gradient changes such that the optimal response (for the prey) exists somewhere between the model and mimic (as Holmgren and Enquist, 1999, claim it does) then the model might evolve towards the mimic in its efforts to get the optimal response from the predator. The authors also explain why the model does not escape the mimic in terms of phenotypic boundaries: a model can be chased to an edge where it has less chance of escape. [3] This is the first neural network model of mimicry, and the approach has the potential to be taken forward. However, neural networks are commonly referred to as “black boxes” as it is often difficult to understand why an evolved network might act in a certain way. Thus, they might prove useful for studies of mimicry and warning signals—but caution should be employed with their use. A notable extension of the model would be to equate one or more phenotypic dimensions with a general level of conspicuousness, in order to make the evolution of warning colouration possible.

6.6.3 Models of Müllerian Mimicry

Müllerian mimicry was not only the first type of mimicry to be modelled (Müller, 1879), but Müller’s work was the earliest mathematical treatment of any evolutionary topic in a Darwinian framework. Müller (1879) suggested that Müllerian mimicry acts to share the burden of predation between the mimetic species. Let a_1 and a_2 represent the numbers of two different defended species. Let n represent the number of individuals killed during the course of a summer while predators are naïve as to its unpalatability. If both species are dissimilar, then each loses n individuals. However, if both species are perfect mimics then the first loses:

$$\frac{a_1 n}{(a_1 + a_2)} \tag{6.2}$$

and the second:

$$\frac{a_2 n}{(a_1 + a_2)} \tag{6.3}$$

---

[3] However, these edge effects may be artificial and their biological plausibility is questionable. Until there is evidence otherwise it would be best to avoid edge effects (see also Bullock, 1999, in this regard).

Thus the absolute gain, due to mimicry, by the first species is:

$$n - \frac{a_1 n}{(a_1 + a_2)} = \frac{a_2 n}{(a_1 + a_2)} \tag{6.4}$$

and for the second:

$$n - \frac{a_2 n}{(a_1 + a_2)} = \frac{a_1 n}{(a_1 + a_2)} \tag{6.5}$$

The absolute gain compared with the occurrence of the species gives:

$$g_1 = \frac{a_2 n}{a_1(a_1 + a_2)} \tag{6.6}$$

and for the second:

$$g_2 = \frac{a_1 n}{a_2(a_1 + a_2)} \tag{6.7}$$

this gives the proportion:

$$g_1/g_2 = a_2^2/a_1^2 \tag{6.8}$$

Recall that naive predators must sample n prey items of a given morphology before learning an aversion. Thus, if defended species share a morphology, then less individuals from each species are sacrificed during the predator's education. The model shows how a resemblance between two prey species reduces the number of individuals sacrificed to sampling or naïve predators during their learning phase. With this work, Müller (1879) sparked interest in using models to study evolutionary phenomena such as mimicry. For derivative or similar models see (Blakiston & Alexander, 1884; Marshall, 1908; Huheey, 1976).

In a seminal paper of a more recent era, Turner (1984) presented a model in support of an argument against a previous and often cited model of Müllerian mimicry developed by Huheey (1976) which used an analytic approach, predicting that Müllerian mimicry, in the conventional sense, should never happen in realistic situations. Huheey (1976) argued that mimicry between two unpalatable species is only ever possible when both species are equally palatable; otherwise the mimicry is always parasitic (as in Batesian mimicry). However, instead of assuming a time element, Huheey's model assumed that predators attacked prey after encountering a given number of prey. As (Turner, 1984, p. 145) puts it, "…on the whole I think it very unlikely that a predator who has decided to give Monarch butterflies a miss, starts counting the number it sees and then attacks, say, the eleventh one! Time-dependent reversal of forgetting seems much more likely to be the general rule.” Turner’s model used a Monte Carlo approach in a similar fashion to the model described in the previous section, although Müllerian and not Batesian mimicry was studied on this occasion. Turner showed that when predator forgetting was modelled as time-dependent instead of encounter-dependent then we should still expect conventional Müllerian mimicry in most situations (although the work on quasi-Batesian mimicry described in section 6.5.2 shows that the matter is not completely settled).

Other notable models of Müllerian mimicry include a population-dynamic study (Hadeler, Demottoni, & Tesei, 1982) and a spatially-explicit (reaction-diffusion) model of mimetic polymorphism (Sasaki, Kawaguchi, & Yoshimori, 2002). Gavrilets and Hastings (1998) examined the dynamics of mimicry evolution using a linear population-genetic approach and found non-equilibruim dynamics, thus suggesting that co-mimics in nature may be in dynamic fluctuation. But it is fair to say that given Müllerian mimicry’s presence right at the birth of the Darwinian paradigm, there are surprisingly few theoretical models of the phenomenon.

6.6.4 Predator Hunger and Nutritional Requirements

There is considerable evidence that predators will attack defended prey as a last resort when they are in severe nutritional need (Poulton, 1890; Swynnerton, 1915; Sexton, Hoger, & Ortleb, 1966; Chai, 1986; Hileman, Brodie, & Formanowicz, 1995; Gillette, Huang, Hatcher, & Moroz, 2000). This suggests that consideration of predators’ nutritional requirements may produce interesting results. Many models of mimicry use a conditioning approach, in which a predator is conditioned by some learning process to avoid certain prey and to attack others (see e.g., Mallet & Joron, 1999). Turner and Speed (1999a) suggested that in addition to this psychological approach it is worth considering predator physiology. They suggest a *saturation* approach, where predators can tolerate the ingestion of a toxin until the concentration reaches some saturation point. Using verbal arguments the authors explain the implication that two prey species that share the same toxin might share a mutually protective relationship even if they are distinct in appearance. This *toxic mutualism* might exist because two prey species that share the same toxin would saturate a predator more rapidly when working together than when working as an individual species.

Holling (1965) is an early model that allows for predator hunger and other such models have recently appeared (Kokko et al., 2003; Sherratt, 2003; Sherratt, Speed, & Ruxton, 2004). Sherratt et al. (2004), for example, used a dynamic programming approach to investigate the effect of shared toxins and predator hunger on predation rates. They showed that when predators are hungry and prey are in short supply it becomes profitable for them to attack mildly defended prey. This results in selection pressure on mildly defended prey to resemble highly defended prey in a relationship that can be either parasitic or mutualistic, depending on the density of the mimetic population. Sherratt et al. (2004) also showed that a species containing one toxin can be a Batesian mimic of a species containing two toxins. This happens because the second species has greater protection in being able to saturate the predator in two ways. Thus, interesting results can be found by considering predator dietary needs. There are numerous ways to take such a model forward. One question that could be asked is "should toxic mutualists be expected to evolve mutualistic mimetic colourations?" (Turner & Speed, 1999a). It might also be useful to incorporate theories from the dietary balancing literature. One such example is the *total minimal discomfort* theory, where animals attempt to minimise their discomfort by balancing their intakes of different nutrients (Forbes, 2001).

6.6.5 The Relationship Between Warning Signals and Mimicry

Although mimicry and warning signals have an obvious relationship, in that a model species must exhibit a warning signal such that a mimic can profitably copy it, thus far all theories and models have considered the two phenomena separately. If we stand back from the separate theoretical debates in each area, we can see how the theories of the origin and evolution of warning signals and of mimicry might interact. From this viewpoint we might make progress on a unified theory. For instance, in our consideration of why so many species exhibit warning signals it could be useful to bring mimicry theory to the table. Consider, for example, a species for which the cost of maintaining a defence might be too high to bear given that the species is not also advertising the defence with a warning signal. If this species could become a Batesian mimic of an aposematic species, it might then be able to "afford" the evolution of a defence. Thus, by considering the two topics together we have both outlined a novel theoretical route for the evolution of Müllerian mimicry and in doing so possibly helped to explain why so many species manage to evolve costly warning signals. In this regard it should be noted that, just as the mimicry of conspicuous colourations is possible, so too is the mimicry of cryptic prey (see Joron, 2003; more will be said on this topic in Chapter 8).

6.7 Summary

Along with warning signals, mimicry is one of the best-known anti-predator adaptations. However, we have seen that many models of Batesian mimicry, and particularly models of Müllerian mimicry, are similar in basic structure and exhibit only slight variations on a theme. A possible reason for this is that current techniques used in the area, such as monte-carlo or population genetic models, struggle to adequately represent the full co-evolution of predators and prey, which is the next natural progression in the development of such models. Thus, theoretical research in the area has been moving slowly. Evolutionary simulation modelling provides a timely and useful tools for advancing models, and consqeuently theories, of mimicry.

Chapter 7

Batesian Mimics Influence the Evolution of Mimicry Rings

Mathematical models of mimicry typically involve artificial prey species with fixed colourations or appearances; this enables comparison of predation rates to demonstrate the level of protection a mimic might be afforded. Fruitful theoretical results have been produced using this method, but it is also useful to examine the possible evolutionary dynamics and consequences of mimicry. To that end, we present individual-based evolutionary simulation models where prey colourations are free to evolve. The models are used to examine the effect of Batesian mimics on Müllerian mimics and mimicry rings (complexes of Müllerian mimics). Results show that Batesian mimics can potentially incite Müllerian mimicry relationships, and encourage mimicry ring convergence.

7.1 Introduction

Mimicry rings are Müllerian relationships between two or more species. Plowright and Owen (1980) showed that there are five different patterns of bumble bee in north-west Europe, each constituting a mimicry ring of several species. Another example of a mimicry ring is the tiger pattern shared by different species of *Heliconius* butterfly (Mallet and Gilbert, 1995, see also, Joron, Wynee, Lamas and Mallet, 2001) which coexists alongside other ring patterns (see, e.g., Figure 7.1). In discussing the formation of mimicry

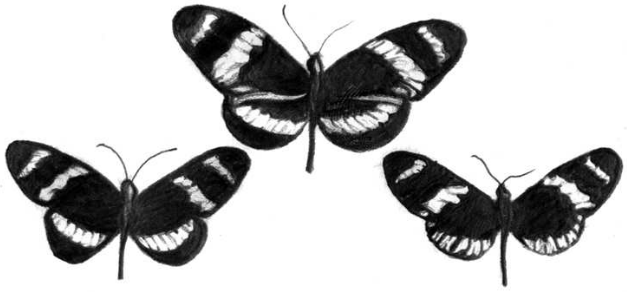

Figure 7.1: Mimicry ring example. Illustration of three separate species of butterfly that all share the same pattern in a ‘yellow’ ring (adapted from Mallet & Gilbert, 1995). Note similarities, and subtle differences, in the patterns.

rings, Turner (1984) makes the analogy to the way planets form. “[L]ike planets forming from a cloud of gas, clusters of mimetic species will arise, and form what we call mimicry ‘rings’. Species occupying spaces in between the rings will be pulled into them, but sooner or later these focal patterns, having absorbed all the available species, will stabilize. If they differ too much from each other they will not be able to converge, for birds will never mistake one for the other.”

It would presumably be highly profitable for a palatable species to invade these rings as a Batesian mimic, and some such relationships are known to exist. The coexistence of multiple mimicry rings within a geographical region raises a question as to why they do not all converge into a single ring for maximum defence from predators (Turner, 1984; Mallet & Gilbert, 1995). Some current explanations suggest that some rings may be seperated due to (subtle) differences in flight height (Papageorgis, 1975; Beccaloni, 1997), differences nocturnal roosting heights (Burd, 1994; Mallet & Gilbert, 1995; Mallet & Joron, 1999) and differences in degree of temporal separation in flight activity (DeVries, Lande, & Murray, 1999) between the animals in the different rings. Ruxton et al. (2004a) point out that “…the evidence is rather equivocal for taxonomic groups such as *Heliconius* (Mallet & Gilbert, 1995).” This question—of why *multiple* mimicry rings coexist—will be looked at briefly before moving on to examine how Batesian mimics might influence the formation of such mimicry complexes.

Mathematical models of mimicry have been useful theoretical tools for examining the costs and benefits to prey as a result of mimicry. They are usually used to examine predation levels on various artificial prey of *fixed* colourations and thus to sketch out the likely evolution of mimetic relationships. However, the coevolutionary dynamics of, and interactions between, different types of mimetic prey have been little explored in models (although see, Gavrilets and Hastings, 1998; Holmgren and Enquist, 1999). Further, it has not been feasible to model the evolution of mimicry rings and the complex multiple predator and multple prey interactions easily using such modelling approaches. In order to address the absence of work on the coevolutionary outcomes of the interactions between different types of mimicry, we present individual-based evolutionary simulation models—in which prey colourations can coevolve—to explore the effect of Batesian mimics on Müllerian mimicry relationships and mimicry rings.

When developing simulations, it is useful to first work with a simpler model to which complex features can be added incrementally. Thus, an initial model is presented in which Batesian and Müllerian mimicry are free to evolve. Thereafter, a richer model is presented to examine the effect of Batesian mimics on mimicry ring evolution to examine how Batesian and Müllerian mimics interact in the formation of mimicry rings.

7.2 A Simple Model of Mimicry

7.2.1 Methods

7.2.1.1 Artificial Prey

Multiple populations of prey species were used in each model variation (which we shall refer to as *experiments*). Different species of prey were each assigned a fixed defence level. Each individual had a single gene: a value representing their phenotype or coloration. The colorations were constrained to a toroidal loop of values from 1–20 (where 20 loops back to 1 and *vice versa*). The values are in a ring in order to avoid edge effects (see chapter 3). The distance of one coloration from another represented their level of similarity.

7.2.1.2 An Artificial Predator

A single predator was modelled with a simple reinforcement learning system. The predator's experience of each coloration was represented by an attack probability; after consuming a prey item of a particular external coloration, the attack probability would be updated according to the defence level of the prey item. The predator generalized on the basis of experience and thus would also, to a lesser extent, update its scores for the closest neighbour colorations. The generalization in the simple model was set to a threshold with a mean window value of size five, with two locations either side of the true value (a value chosen to create a landscape of similar, less similar, and distinct prey phenotypes). The attack probability gradually degraded back towards ambivalence at a constant rate of two percent per prey generation.

7.2.1.3 Model Conditions

For initial simplicity, in each generation the predator was presented with 30 binary forced-choice situations. Two individuals were randomly selected from across all prey populations present and the predator would make a probabilistic attack decision, based on its experience of each coloration. Random reproduction then took place amongst the surviving prey. Mutation was implemented as a uniform change of ± 1 in the coloration, and the mutation rate was 0.03. All of the experiments were run over 5000 generations, and prey species populations kept constant at 100. The main variable manipulated was the starting distance between prey species' colorations, in order to determine whether an initial chance resemblance is required for the evolution of mimicry.

7.2.2 Results

Figure 7.2 shows the results of two experiments in terms of the initial and final distances between prey species' colorations. Experiment 1 was conducted using one palatable and one defended species. Regardless of the starting distance between the two, Batesian mimicry evolved: the palatable species came to have the same or very similar coloration to the defended one. This was also true when the function to update the predator's memory did not include generalization.

Experiment 2 was conducted in the same way with two defended species. Figure 7.2 shows that Müllerian mimicry only evolves if the two prey species have some initial resemblance. If they are initially more than approximately four units away on the phenotypic ring, they typically remain distinct in appearance. Selective pressure dictates that defended species require an initial resemblance to promote a Müllerian relationship.

In experiment 3 (no graph shown) simultaneous Müllerian and Batesian mimicry was investigated by including two defended and one palatable species. The results showed that the coloration of the palatable species moved towards that of one of the defended species (i.e., Batesian mimicry). Interestingly, this in turn drove the model species around the phenotypic ring and resulted in Müllerian mimicry with respect to the second defended species, regardless of their initial phenotypic distances (when plotted the results look similar to those of experiment 1).

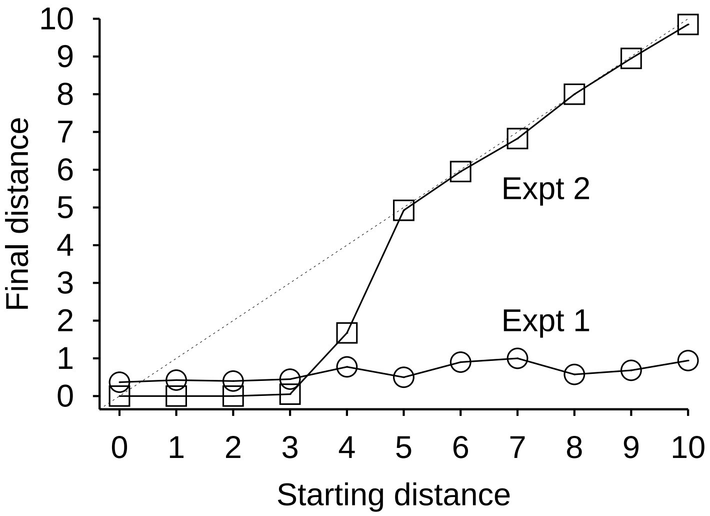

Figure 7.2: Final distance by initial distance between two prey species' colorations; each point is averaged over 40 runs. Dashed line shows zero change in coloration.

7.2.3 Discussion

Experiment 1 shows that Batesian mimics will close in on the model regardless of how different their initial phenotypes are, and regardless of the predator's ability to generalise. This shows that there is more selective pressure for the palatable species to resemble the model than there is for the model to diverge. However, there is an additional reason why the mimic is successful in 'catching' the model: before mimicry has evolved, palatable individuals gain an inherent fitness benefit for mutating away from the average (modal) phenotype of their species. This is because, when presented with a choice, the predator would be more likely to select a well-known phenotype than a newer one.

The results of experiment 2 show that Müllerian mimicry relies on an initial resemblance between the species involved. Such initial resemblance might be caused by any number of factors, such as sexual selection, random drift, phylogenetic similarity, and so forth (it should be noted that the *particular* resemblance threshold found in the experiment—in this case four units on the phenotypic ring—is of course determined by the nature of the predator's generalization.)

Experiment 3 demonstrates an important finding: that pressure due to Batesian mimicry can be a force that drives Müllerian mimics together despite a lack of initial resemblance.

7.3 A Model of Mimicry Ring Evolution

The above results suggest that Batesian mimics can influence Müllerian mimicry relationships. This implies that Batesian mimics may influence the formation of mimicry rings. In the previous simulation the prey colorations were represented only by a single gene. A richer model will now be presented which incorporates multiple dimensions of prey coloration, and explicitly examines mimicry ring evolution.

7.3.1 Methods

7.3.1.1 Artificial Prey

Artificial prey were modelled as populations of individuals with a fixed defence level and a genome representing their coloration. Multiple populations of prey were used in each experiment, and each was assigned a palatability level on a scale between zero and one (most to least defended respectively). The genome consisted of multiple genes with values compositely representing prey coloration; most of the results presented here refer to the two-gene case. The genes were constrained to *continuous* values from 1–200 and the Euclidean distance of one coloration from another represented the degree of similarity between phenotypes.

7.3.1.2 Artificial Predators

A population of predatory individuals was modelled with a Monte Carlo (Hebbian style) reinforcement learning system, as used in a mimicry model by Turner et al. (1984). Predators were presented with ten prey per prey generation, which they could either consume or reject. Through experience, the simulated predators could learn to associate prey colorations with a level of defence. As such, they could adapt their probabilities for each coloration as they experienced them, according to the defence level of the consumed prey. Turner *et al.*'s (1984) chosen reinforcement learning system was modified to include generalization, in order to capture the notion that predators will treat similar colorations in a similar way. Predators were modelled as having an attack rating (see below) for each prey coloration. The formula used for generalizing and *updating* these attack ratings after consuming a prey item was:

$$P_{t+1} = W(P_t + \alpha(\lambda - P_t)) \tag{7.1}$$

This produces an updated attack rating P_{t+1}, based on the rating at the previous time step (P_t), for colorations at a given similarity distance from the consumed prey whose defence level was λ. The α term denotes a variable learning rate and is calculated using:

$$\alpha = 0.5 + |\lambda - 0.5| \tag{7.2}$$

W is a weighting (used for generalization) which is calculated according to the distance of a coloration from the consumed prey's coloration. If $D > G$ then $W = 0$ (i.e., generalization is finite in extent) otherwise:

$$W = \frac{G - D}{G} \tag{7.3}$$

Where G is the predator's generalization range and D is the Euclidean distance between the coloration and the consumed prey's coloration.

The learning rate is dependent on how far predators' current attack ratings are from the true defence level of the attacked prey (Equations 1 and 2). Predators' memories are updated by generalizing over similar colorations within their generalization range, moving outward from the coloration of the attacked prey item (Equation 3). This update can be visualized as a cone shape, where the cone tip represents the update for the consumed coloration. The further the coloration from the one consumed, the less it is updated.

Attack ratings were transformed with the logistic function to give a probability of attack, such that predators were decisive about prey for which they had a relatively strong opinion. Whereas in the previous simulation predators were forced to consume at least one prey item per prey offering, here a more realistic assumption was implemented, where attack decisions were made by comparing the attack probability to a pseudo-random number in the range of zero to one. Thus, it was possible for predators to become averse to eating some or even all prey species. We therefore assume that an additional food source is available; predators in the model will not starve even if they refuse all prey items. Predators' memory degraded over time by gradually reverting back toward ambivalence (i.e., an attack probability of 0.5) at a constant rate of two percent per prey offering (as used by Turner *et al.*, 1984). Random asexual reproduction then took place among the surviving prey. After every prey generation the oldest predator would die and be replaced with a new and naïve predator.

7.3.1.3 Model Conditions

All of the conditions were run over 200,000 generations and prey populations were kept constant after reproduction. Predator generalization (G) was set to a value of eight (this value should be considered relative to the abstract phenotypic space which was 200 units across in each dimension). Stepwise mutation (choosing a random gene to mutate) is commonly used in simulation models, but it has an inherent bias towards orthogonal directions in multidimensional genetic space. This could influence the results of the simulation. To overcome this problem a mutation operator was implemented as follows:

1. A random direction in the multidimensional space was chosen by selecting a random number from a normal distribution (0 mean, unit variance) for each gene.
2. The distance was then selected over a normal distribution (0 mean, unit variance) to allow for varying mutation sizes, with a bias towards smaller ones.
3. The offspring were then mutated in the selected direction at the selected distance.

Note that the mutation operator took place on *every* offspring. The result is that most, but not all, offspring are *slight* variants of their parent with the possibility of rare extreme variants. The issue of mutation bias due to unnatural boundaries (Bullock, 1999) was handled using a wrap-around function (making the values 0.0 and 200.0 synonymous) on prey mutations and predator generalization. Attending to such implementation issues is important for avoiding artefactual results in models. They are of particular significance to mimicry models as random drift and mutation are vital to the initiation of mimicry. We consider this type of mutation operator appropriate for this type of simulation, in which there is no distinction between genotype and phenotype, and where the prey representations are additive-genetic in character.

Clustering algorithms were considered as a method for assessign the number of mimicry rings formed at the end of each simulation run. However, cluster algorithms are quite time consuming and did not give results that matched our manual observations well. For our purposes, a mimetic relationship is defined as a Euclidean distance of less than fifteen between the modes of two species. The distance of fifteen was an approximation of how close species needed to be to be classed as being in a ring [1] Although this is, of course, not a standard cluster analysis techniue, it was used because it is not processor intensive, and it gaves clusters matching those found through manual observation of the results.

7.3.1.4 Experiment

We ran the simulation with various numbers of palatable prey (i.e., potential Batesian mimics). For the results given below, the prey coloration was represented by two genes.

---

[1] Calculated by multiplying the generalization range by two and deducting one (which gives the minimum distance of overlapping generalization) and then multiplying that by 1.5 (to allow for spread of individuals within a species and generalization over each of them).

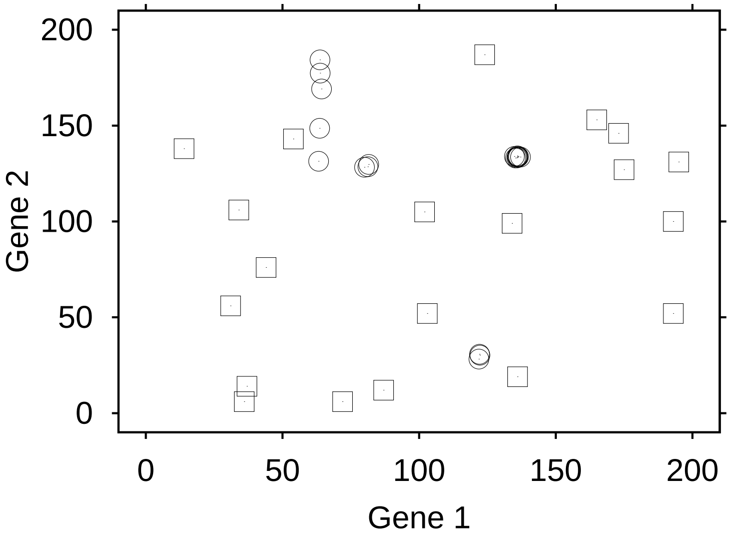

Figure 7.3: Initial random and final evolved position of each prey species' modal coloration for a typical run with equal population numbers. Squares represent starting colorations and circles represent final colorations.

A random initial colour pattern was chosen for each prey species at the start of each run. A defence level of 0.9 and 0.1 was given to defended and palatable prey respectively. All prey population sizes were kept constant at 300. The simulation results for each experiment were recorded over 20 runs, after which the number of coexisting mimicry rings were tallied.

7.3.1.5 Sensitivity Analysis

The same experiments were also performed with four genes (and an equivalent mutation rate and range of genetic values giving the same sized space of colorations as the two gene case, using the generalization radius as a reference point). The four-gene results were not significantly different from those presented here. The two-gene results are representative of our findings and are presented here for ease of conveying and visualizing the results. A sensitivity analysis (systematically varying the values of some key parameters) showed that the results are robust over a range of values. The parameter space tested was: dimensions of prey phenotype [1,2,3,4], predator generalization [4-14], prey population size [200-600], defence level of undefended prey [0.7-1.0].

7.3.2 Results

Figure 7.3 shows the results of a single representative run with no palatable species present. Figure 7.4 shows the results for a run when four palatable prey species were

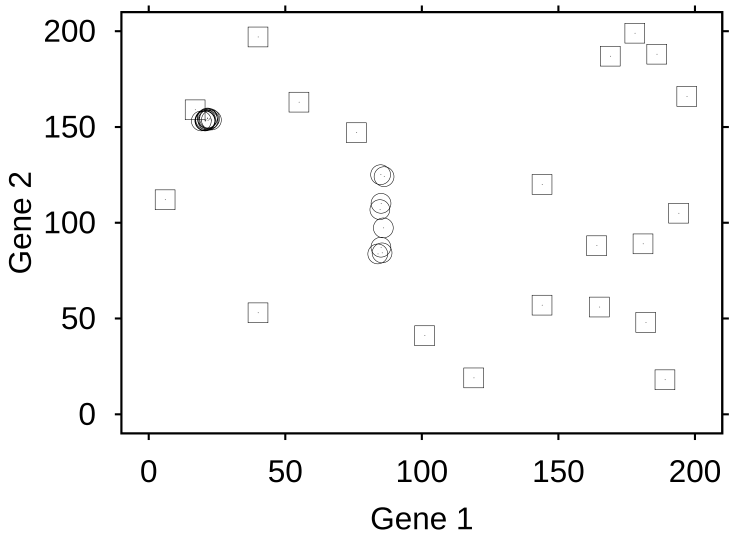

Figure 7.4: Initial random and final evolved position of each prey's modal coloration for a typical run with four palatable species added. Squares represent starting coloration and circles represent final colorations. Palatable species are not shown; defended species form two large rings and three small rings. Notice how there are fewer rings when palatable species are present.

present. The important points to take from these two figures are that mimicry rings form in both cases (i.e., the final colorations are clustered) and the number of rings or clusters is smaller in Figure 7.4. The latter result indicates that Batesian mimics have encouraged the formation of a smaller number of large mimicry rings.

Figure 7.5 shows that as the number of palatable prey species present increases up to three, the number of mimicry rings decreases. Adding more than four palatable species typically results in a number of mimicry rings not significantly different from adding just four palatable species.

A statistical comparison of the number of mimicry rings found in each of the two experiments indicated a significant difference ($t = 5.07, p < 0.001$) with more rings being found with no palatable prey added than with four.

7.3.3 Discussion

When no palatable species were present in the model ecosystem (i.e., there was no possibility of Batesian mimicry) multiple mimicry rings evolved. A further examination of the results showed that in all runs there is little change from generation 150,000 to 200,000. This long-term stability occurred *despite* the fact that mutation rates were high (every offspring was a mutant) and variation in prey coloration was cost free. If multiple mimicry rings coexist and do not merge into a single ring despite these conditions, we can see that individual rings, once formed, are capable of maintaining their identity in the face of perturbations due to random drift. Of course, if predators generalized over the entire phenotypic space we would expect all prey to evolve into a single large ring, but under more realistic circumstances presented here, multiple mimicry rings evolve. The number of mimicry rings would obviously fall as the generalization range is increased. However, we have shown that given the assumption that when predators do not generalize to a great extend (e.g., they do not generalize between blue and red) then distinct mimicry rings can form. The results complement explanations for the diversity of mimicry rings in nature (see e.g., Turner , 1977; Sheppard *et al.*, 1985) which suggest that the colorations of different mimicry rings may be so dissimilar that the rate of major mutations that could allow a species to swap rings tends towards zero. This is because intermediate mutational forms are at a selective disadvantage.

Introducing palatable species to the model shows that Batesian mimics can influence the evolutionary convergence of mimicry rings. In the model, palatable species inevitably become Batesian mimics and chase their respective mimicry rings through cycles of colorations, increasing the chance that two mimicry rings might move within convergence range of each other. There is also a threshold over which an increase in the number of palatable species does not *further* reduce the number of mimicry rings. This is because, although Batesian mimics can entice Müllerian relationships, an excess of Batesian mimics can break up or prevent them (Pough et al., 1973). We verified this by running the simple simulation with an excess of Batesian mimics. The result was that Müllerian mimicry almost never occurred with respect to two defended species.

7.4 A Mathematical Model of Mimicry Rings

Ruxton et al. (2004a) (in a book published *after* the publication of the model presented in this chapter) provided a mathematical treatment of mimicry rings to look at the question of why multiple mimicry rings exist. They looked at the evolution of 20 equally unpalatable species whose phenotypes are represented in two dimensions: x and y. Each species begins with a random value for x and y, selected from a random distribution within the limits of 0 and 10. The phenotype distributions of each species were assumed (in contrast to our approach) to be normal with standard deviation δ. Each time-step a mutant was created in a randomly selected species. The mutant prey item differs from the modal phenotype magnitude by δ_x and δ_y respectively—both increments were selected from a normal distribution (mean 0, standard deviation 1). Predators acted to maximize benefits

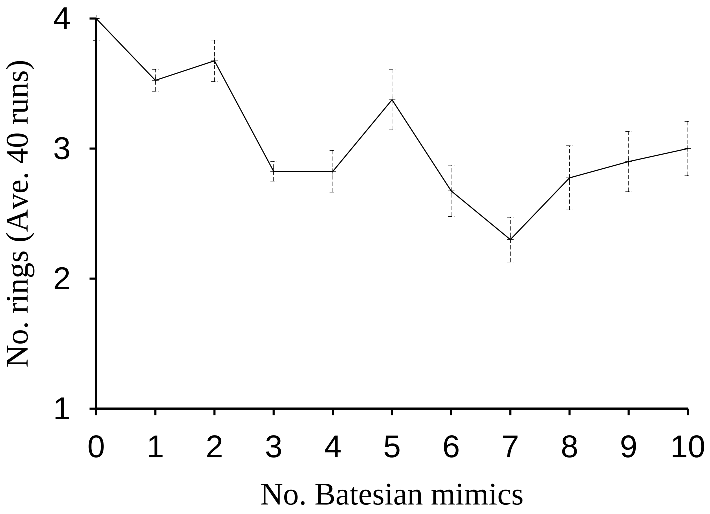

Figure 7.5: The mean number of coexisting mimicry rings formed when a various number of palatable (potential Batesian mimics) prey populations are present.

of attacking undefended prey and mimimize costs of attacking defended prey. [2] Ruxton et al. (2004a) assumed that mutation events were so rare that if they occurred in any species, that the mutation would either fixate within the population or become extinct. The authors justified this "invasion implies fixation" rule on the basis that as the density of mutants increases, so does their per capita survivorship. Although this is actually more of an assumption than a justification and may restrict drift, it is a reasonable assumption for the purpose of the particular model. The model was run under different conditions where the species were either equally abundant or where one species has a density much greater (100 times) than the other species. Their results (see Figure 7.6) complemented the initial result of the simulation presented in this chapter—distinct mimicry rings evolved because intermediate mutant forms are at a selective disadvantage. In addition, when species with a high population density were introduced to the model they acted as hubs: "When one [undefended] species is particularly common in the system then, as expected, its own appearance did not change over time and other [undefended] species evolved towards it." Ruxton et al. (2004a) concluded that evolutionary pathway taken by mimicry rings depends on the initial conditions and the limited probability of large mutations. Early trial runs with the evolutionary simulation model presented in this chapter were consistent with these results: dense defended prey populations could act as hubs for mimicry rings.

---

[2] The exact way in which predators were modelled is not documented.

7.5 Conclusion

Our model has shown that multiple Müllerian mimicry rings might exist for exactly the reasons that Turner (1984) suggested with his planet formation analogy: that intermediate mutational forms (mutant prey with phenotypes that do not closely resemble that of any ring) are at a selective disadvantage.

The results of evolving mimicry rings with no possibility for Batesian mimicry showed that, under conditions in which all species are (equally) defended, multiple mimicry rings evolved. When we consider this, along with the high mutation rate (every offspring is a mutant), cost-free variability of prey coloration, and so forth, we can see that ring stability, due to convergence, could be a major factor in the diversity of mimicry rings. A further examination of the results showed that in all runs there is little change from generation 150,000 to 200,000. The results complement suggested explanations for the diversity of mimicry rings (see, e.g., Turner , 1977; Sheppard *et al.*, 1985, which suggest that the colourations of different mimicry rings may be so dissimilar that it takes a long time waiting for a large major mutation before it is possible for a species to swap rings.

The results of introducing palatable species show that potential Batesian species can influence the evolutionary convergence of mimicry rings. In the model, palatable species inevitably become Batesian mimics, effectively destabilize the rings and 'chase' the rings through cycles of colorations, increasing the chance that mimicry rings would arrive within convergence range. There is also a threshold over which an increase in the number of palatable species does not *further* reduce the number of mimicry rings. This is because, although Batesian mimics can entice Müllerian relationships, an excess of Batesian mimics can break-up or prevent them. We confirmed this explanation by running the simple simulation with an excess of Batesian mimics. The result was that Müllerian mimicry was rarely sustained with respect to two defended species.

The results of both simulations suggest that Batesian mimics can promote Müllerian relationships and, furthermore, influence mimicry ring evolution. This finding has empirical consequences. If Batesian mimics drive defended species towards Müllerian mimicry, it follows that larger Müllerian mimicry rings should be expected to have a proportionally greater number of Batesian parasites than will smaller rings. Experimental tests for these predictions are intended to take place using computerized prey and human predators (Tom Sherratt, personal communication).

Gavrilets and Hastings (1998) have examined dynamics between Batesian and Müllerian mimics previously with an elegant mathematical model, but the individual-based mimicry ring model presented here is the first to look at interactions between *multiple* potential

Batesian and Müllerian species. Ruxton et al. (2004a) state that the model presented here "is important because it marks a shift in emphasis away from a simple dichotomous view of two unpalatable prey and a single predator to a more realistic community perspective involving multiple Batesian and Müllerian mimics."

Finally, we would echo a question raised by John Turner (personal communication). Given that Batesian mimicry occurs so readily in both of the simulations presented, why is it not more common in nature? This question appears simple on the surface, but is actually loaded with substance. No mimicry model to date allows evolution along a dimension related to conspicuousness—and the model presented here follows suit. This is not a problem for this particular model, as we have chosen to examine the effect of Batesian mimics on mimicry rings and not the tradeoff between mimicry and crypsis. But not considering conspicuousness and its attendant costs is a possibly important omission for models of mimicry in general. What happens if we build a model that has room for both conspicuous warning signals and mimicry? This leads us conveniently on to the next chapter.

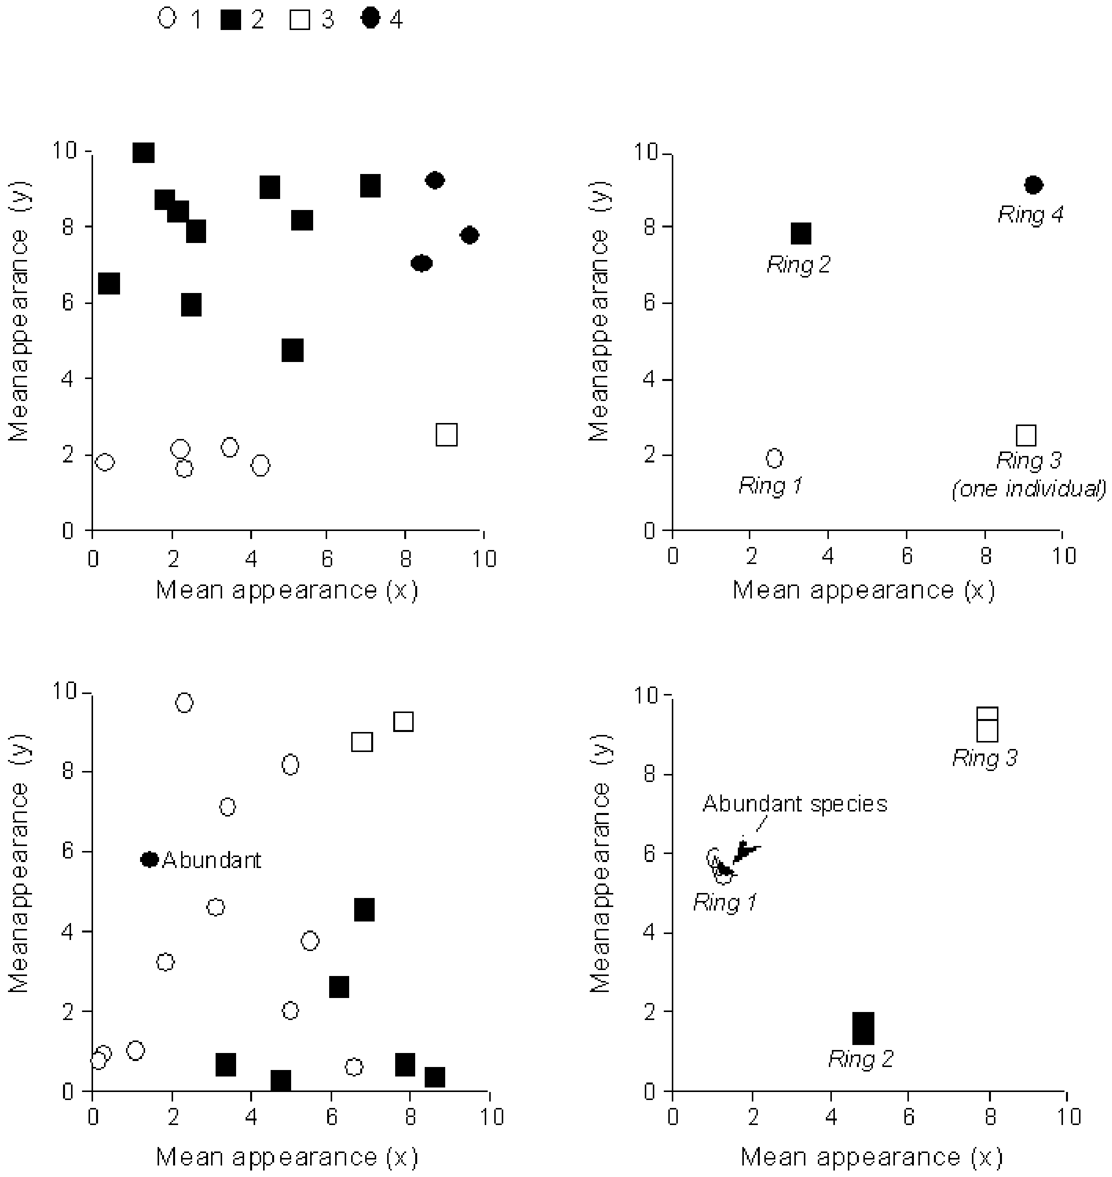

Figure 7.6: Mimicry ring formation among 20 unpalatable species. (a) and (c) show the initial conditions (randomly selected) in two separate simulations. (b) and (d) show the respective outcomes from these two simulations after 10000 iterations. The different mimicry rings are assigned a shape (such as filled circle, empty square, etc.) in (b) and (d), and the species in (a) and (c) are displayed with the shape assigned to the mimicry ring to which they were eventually drawn. Graphs taken from (Ruxton et al., 2004a).

Chapter 8

Warning Signals as a Response to Mimicry

8.1 Introduction

Up to now, this thesis has, for the most part, discussed theories of warning signals and mimicry individually. This echoes the way in which research in the area has typically been carried out in the past—despite the obvious relationship between the two phenomena. How do warning signals and mimicry relate? This chapter represents a first attempt at consolidating warning signal and mimicry theory (see, Sherratt & Franks, 2004). In doing so, a contribution is also made to the issue of whether warning signals evolved solely to exploit preexisting predator biases, or to distinguish defended from undefended prey.

To recap, undefended prey tend to be cryptic and defended prey tend to be conspicuous. For example, defended prey tend be bright, aggregate, and move slowly (Bates, 1862). There is even evidence that some defended prey might be willing to pay an energetic cost for slow flight (Srygley, 2004). Returning to a question posed in chapter 4: why have these conspicuous traits been selected for in defended prey? One possibility is that defended prey are taking advantage of opportunities available to conspicuous prey. Conspicuous prey are able to roam openly; as such prey are already conspicuous, to some extent, they can take advantage of benefits such as easier foraging, more elaborate sexual displays, and thermoregulation. As Turner (1984) points out, warning colourations may allow a species to come into the open and "… operate on solar power, flying by day being much more economical of energy than flying by night." (e.g., Wasserthal, 1975; Douwes, 1976; Douglas, 1979).

This chapter readdresses the question of why conspicuous traits have evolved in conspicuous prey, and asks if defended prey have evolved traits such as conspicuous warning colourations and aggregations in order to escape being parisitised by Batesian mimics. Thus, this chapter presents the first model to combine warning signals and mimicry. Although this thesis is mainly concerned with evolutionary simulation models, the model presented in this chapter is mathematical. However, this is the first model of its kind, and is intended as the first step towards building a detailed evolutionary simulation model. The reader is referred to Chapter 3 for arguments in favour of laying down a mathematical framework before constructing an evolutionary simulation model. Plans for the next step in the modelling process (i.e., constructing an evolutionary simulation model) are presented later in this chapter.

8.2 The Contemporary Theory

In chapter 4 we saw that there are two main stances on the evolution of warning signals—the 'receiver psychology', and 'conspicuousness as a reliable indicator of defence' approach. For more detail on these theories, the reader is referred back to Chapter 4. This chapter briefly returns to this discussion before presenting a relevant model.

It has been repeatedly shown that predators tend to learn to avoid conspicuous defended prey quicker than cryptic defended prey. In addition it has been shown that forgetting is decelerated for conspicuous prey (i.e., predators can remember experiences with conspicuous prey for longer). The conventional 'receiver psychology' theory of warning signal evolution stresses the observations of such educational tendencies of predators, proposing that warning signals evolved to exploit such biases (e.g., Gittleman & Harvey, 1980); that there is "something special" about conspicuousness. This approach was championed by Guilford and Dawkins (1991) and has become one of the most popular theories in the field.[1] Epitomizing this stance (Lindstrom et al., 2001, p. 9181) states that "Aposematic species are conspicuously colored because predators learn faster to avoid conspicuous patterns."

Sherratt and Beatty (2003), insightfully, highlighted an important caveat in the assumptions underlying the receiver psychology approach; whatever the underlying cause of warning signals, predators are likely to have evolved such predispositions in the first

---

[1] Guilford and Dawkins (1991) do, however, brush over discriminative learning.

place in response to warning signals. Thus, the empirical observations of predator educational biases favouring the survivorship of aposematic prey are not *necessarily* evidence that such biases are responsible for the evolution of warning signals. However, such observations can certainly help to explain why warning signals are maintained. Ruxton et al. (2004a) also make this point, asking “what came first, conspicuousness or special predator responses to conspicuousness?” In other words, are the observed predator biases ‘just one of those things’, or have they adapted in response to evolved distributions of defended and undefended prey across cryptic and conspicuous forms. Whether these responses necessarily existed at an early stage in the evolution of aposematism remains an unresolved coevolutionary question that is extremely difficult, if not impossible, to make empirical progress on. Turner (1984) stated that “. . . the coevolution of the vertebrate nervous system with the insects may have made it impossible for us finally to disentangle the problem.” However, models offer us a chance to take such studies forward. We can, at the very least, clarify and explore hypotheses with progressively more complex theoretical models.

8.3 How Might Warning Signals Become Reliable Indicators of Defence?

In Chapter 4 we discussed the theory that warning signals have evolved in defended prey to distinguish them from undefended prey, thus reliably indicating a defence. Contemporary receiver psychology theories of warning signals have tended to focus on associative learning and the fixed educational biases of predators. However, lateral approaches such as those based on peak shift (Yachi & Higashi, 1998; Leimar et al., 1986) and predator filtering effects (Sherratt, 2002a; chapter 5) and the approach here have explored the importance of discriminative learning for the evolution of warning signals.

Early discussions of warning signals tended to centre around the idea that defended prey are under selection to evolve traits that distinguish them from undefended prey. In Chapter 1 we saw how Wallace and Darwin puzzled over why animals are brightly coloured, and how Wallace (1867, 1879b) was the first to provide a possible explanation. In *Darwinism* (Wallace, 1889) he theorised that:

> . . . the animals in question are possessors of some deadly weapons, as stings or poison fangs, or they are uneatable, and are thus so disagreeable to the usual enemies of their kind that they are never attacked when their peculiar powers or properties are known. It is, therefore, important that they should

> not be mistaken for defenceless or eatable species of the same class or order since they might suffer injury, or even death, before their enemies discovered the danger or uselessness of the attack. They require some signal or danger flag which shall serve as a warning to would-be enemies not to attack them, and they have usually obtained this in the form of conspicuous or brilliant colouration, very distinct from the protective tints of the defenceless animals allied to them.

More famously, Fisher (1930), following in Wallace's footsteps, proposed that "selection will tend to modify the model so as to render it different from the mimic and as conspicuous as possible." This theory gradually became overshadowed by the receiver psychology approach (possibly due to the ease with which the latter theory allowed empiricists to design experiments).

We now return to a question posed in Chapter 4. Why do defended prey typically use conspicuous traits—such as bright colourations, slow flight, and aggregations—as signal components to distinguish themselves from undefended prey? Why have they not simply evolved along traits that are unrelated to conspicuousness? One possibility is that defended prey are availing themselves of opportunities afforded to them by conspicuousness. For example, conspicuous prey can freely roam to forage and thermoregulate, benefit from sexual selection, and so forth.

There are different ways to be cryptic. "Taking random samples from the background would suggest that, providing the background is spatially variable, there can be a number of background-matching forms ..." (Ruxton, Sherratt and Speed, 2004, p. 13). Thus, Batesian mimicry is a possibility among cryptic prey. Joron (2003) reasoned that "rampant mimicry", found in a group of inconspicuous (transparent) ithomiine butterflies, demonstrates that predators are able to learn to avoid such prey. Given that defended prey could be easily mimicked when cryptic, conspicuousness might have evolved in defended prey because it is a difficult trait for undefended prey to exploit and sustain. Thus defended prey might be expected to evolve traits that undefended prey find difficult to exploit, in order to shake off parasitic Batesian mimics. Conspicuous colourations, aggregations, slow flight, and so forth, might serve as traits that are difficult for undefended prey to forge.

Although the idea that conspicuousness has evolved in defended prey as a means for them to shake off their Batesian mimics appears intuitive, little has been done to question the validity of the argument, or examine its implications. For example, there are no models of mimicry that explicitly allow for the evolution of conspicuousness, and no models of warning signals that allow for levels of mimicry.

Models of warning signals have typically allowed prey evolution along a level of conspicuousness, with no room for mimicry along traits unrelated to conspicuousness. Conversely, models of mimicry have always modelled the appearance of an animal along traits independent of conspicuousness. It is clear that mimicry and warning signals are highly related phenomena, and this omission in current models could have important implications. To examine the question “Do defended prey evolve traits that undefended prey find difficult to exploit?”, we constructed a model that allows for both conspicuousness and mimicry.

8.4 The Model

To begin, we assume two separate species, defended prey D and undefended prey U, which can each take two separate forms (1 and 2). Thus, the densities of different prey forms are denoted $D1$ and $D2$ for defended, and $U1$ and $U2$ for undefended prey (for a similar approach see Gavrilets and Hastings, 1998). The two parallel forms of each species ($D1$ and $U1$, $D2$ and $U2$) look alike to predators to a given degree (measured as $s1$ and $s2$, $0 < s1, s2 < 1$, see below). However, we assume that forms 1 and 2 are sufficiently distinct that predators do not confuse them. Predators are awarded benefit b for attacking undefended prey, and incur cost c for attacking defended prey. $q1$ and $q2$ represent the probability of a predator detecting forms 1 and 2 respectively; thus we can assume that they are directly proportional to prey conspicuousness. We also assume a per capita rate of mutation m and logistic population growth parameters r and K, which are identical for both species. With these parameters, we give the dynamical rate equations for $D1$, $D2$, $U1$, and $U2$ (the change in the frequency of the distinct forms of each species) as:

Defended prey

$$\frac{dD1}{dt} = rD1[1 - (D1 + D2)/K] - mD2 + mD1 - f_1(b, c, s1, D1, U1)q1D1 \quad (8.1)$$

$$\frac{dD2}{dt} = rD2[1 - (D1 + D2)/K] - mD1 + mD2 - f_2(b, c, s1, D2, U2)q2D2 \quad (8.2)$$

Undefended Prey

$$\frac{dU1}{dt} = rU1[1 - (U1 + U2)/K] - mU2 + mU1 - g_1(b, c, s2, D1, U1)q1U1 \tag{8.3}$$

$$\frac{dU2}{dt} = rU2[1 - (U1 + U2)/K] - mU1 + mU2 - g_2(b, c, s2, D2, U2)q2U2 \tag{8.4}$$

The first part of each equation represents growth, and the middle section mutation in and out. The functions $f_1()$, $f_2()$, $g_1()$, $g_2()$ represent predation, and refer to the probability of attack of a given form on encounter with a predator. Signal detection theory is used here to provide an optimal solution for the predator, based on the relative probabilities that an encountered prey type is defended or undefended, weighted by the costs and benefits of attacking these prey types.

Signal detection theory is a general mathematical principle that can be applied to any problem where there are two stimuli, with noise, that need to be discriminated from one another. In this case, there can be four possible outcomes from predator attacks: hit (predator attacks an undefended prey item), miss (predator avoids an undefended prey item), false-alarm (predator attacks a defended prey item), and correct-rejection (predator avoids a defended prey item). The two stimuli ($D1$ and $U1$ or $D2$ and $U2$) may be dissimilar to a certain degree, but because of noise (due to recognition errors, generalization, and so forth) surrounding each stimuli the signals can be confused. Thus, the closer the resemblance of the two stimuli (higher values of $s1$ and $s2$), the more the stimuli overlap and the more difficult it becomes for the predator to distinguish between them. Signal detection theory is a means for calculating the optimal (criterion) response of a predator given this imperfect information. A tutorial on signal detection theory is beyond the scope of this thesis, but the reader is referred to any of a number of available accounts (Oaten et al., 1975; Staddon & Gendron, 1983; Greenwood, 1986; Getty, 1985; Getty et al., 1987; Sherratt, 2001; Johnstone, 2002; Sherratt, 2002b).

Using Staddon and Gendron's (1983) power curve approximation (see also Greenwood, 1986) for the signal detection gives us:

$$g_1() = [(U1 * b * s1)/(D1 * c)]^{(s1/(1-s1))} \tag{8.5}$$

$$f_1() = g_1()^{(1/s1)} \tag{8.6}$$

Subject to the constraint:

$0 \leq f_1(), f_2(), g_1(), g_2() \leq 1$

This signal detection component allows us to model predator generalization and recognition errors. With parameters $s1$ and $s2$ we can control the level of mimicry between the parallel forms of the two species by changing the frequency that defended and undefended prey of the same form are confused by predators. For example, when $s1 \rightarrow 1$ (perfect mimicry) then the effect on the subsequent behaviour of predators attacking defended and undefended prey is the same $f_1() = g_1()$. Conversely, if $s1 = 0$ (both species of form 1 are distinct) then $g_1 = 1$ while $f_1() \rightarrow 0$. When $s1 > s2$, then predators can more easily discriminate between defended and undefended prey when they take form 2 than when they take form 1.

8.5 Results

Equations 8.1-8.4 were solved numerically from fixed starting points $D1_0, D2_0, U1_0, U2_0$ using a fourth-order Runge-Kutta approach. When population growth was high enough to replace individuals lost due to predation (such as, $r = 1$, $K = 1000$) then the proportion of individuals of a given species in one form or another was insensitive to the precise growth rate parameters. This was confirmed by repeated numerical integration. The mutation rate m was assumed low in all cases ($m = 0.001$), such that predation was the primary source of population-level change for the prey. To explore the implications of the model we first examine extreme cases of perfect mimicry and no mimicry, before exploring more general conditions. In all cases, we are interested in what forms the defended and undefended species evolve to adopt.

1. The Parallel Forms Are Perfectly Mimetic

In this first extreme case both parallel forms are completely mimetic $S1 \rightarrow 1$ and $S2 \rightarrow 1$ (Figure 8.1*a*). In this case one form is slightly more conspicuous than the other ($q2 > q1$). When the cost to a predator of attacking defended prey is lower than the benefit of attacking undefended prey ($c < b$) then undefended prey tend to undermine any conspicuous display that evolves in defended prey. Thus, under these conditions both defended and undefended prey evolve to adopt their most cryptic form (the lowest common denominator), regardless of the starting conditions. In contrast, when $c > b$ then the optimal behaviour for the predator is to avoid all prey, as it cannot distinguish the highly costly

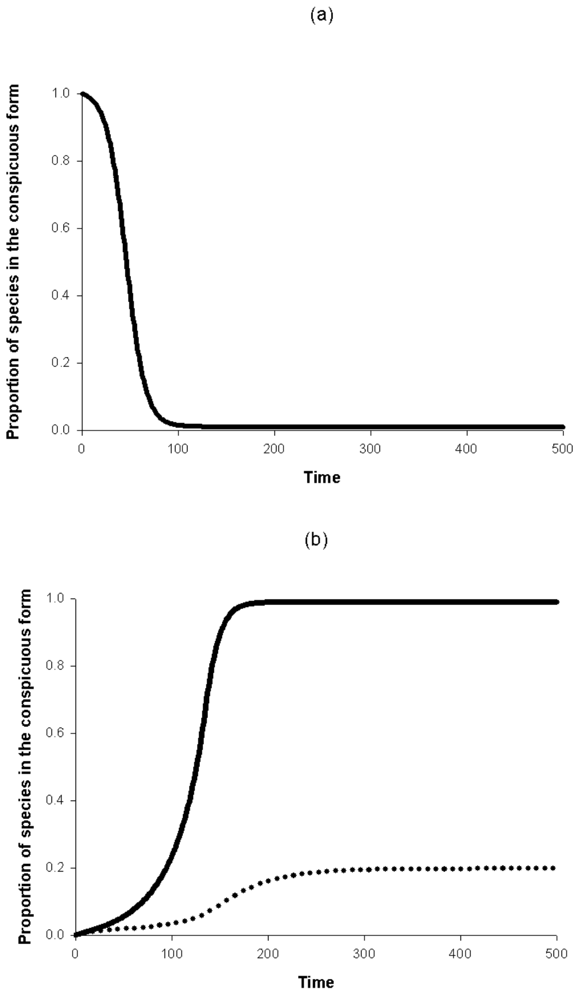

Figure 8.1: Evolution of profitable and unprofitable species' appearance as a consequence of optimal predator behaviour. The solid line denotes unprofitable prey, and the dotted line denotes profitable prey. (The dotted line in (a) is overlaid by the solid line) Parameter values: $b = 2$, $c = 1$, $m = 0.001$, $r = 1$, $K = 1000$, $q1 = 0.1$, $q2 = 0.2$, $D1_0 = 0$, $D2_0 = 500$, $U1_0 = 0$, $U2_0 = 500$. In *1a*: $s1 = 0.999$, $s2 = 0.999$ (perfect mimicry). In *1b*: $s1 = 0.2$, $s2 = 0.2$ (imperfect mimicry).

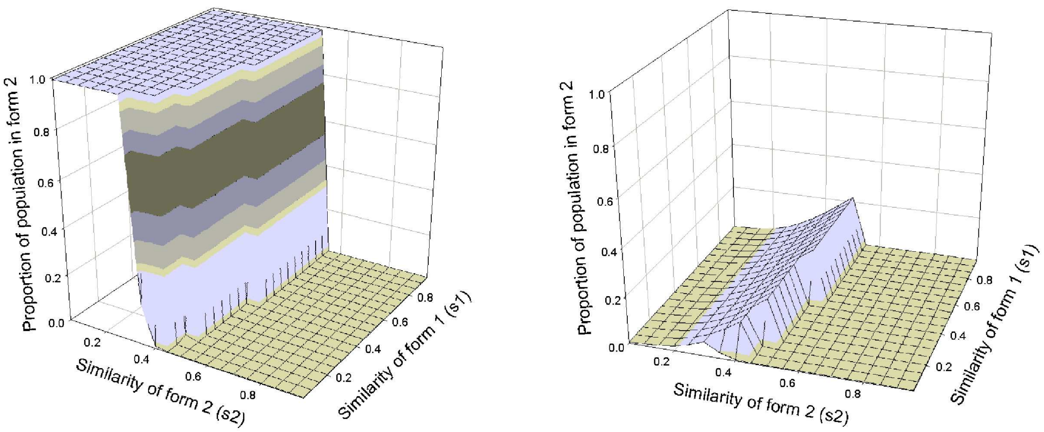

(a) Defended Species

(b) Undefended Species

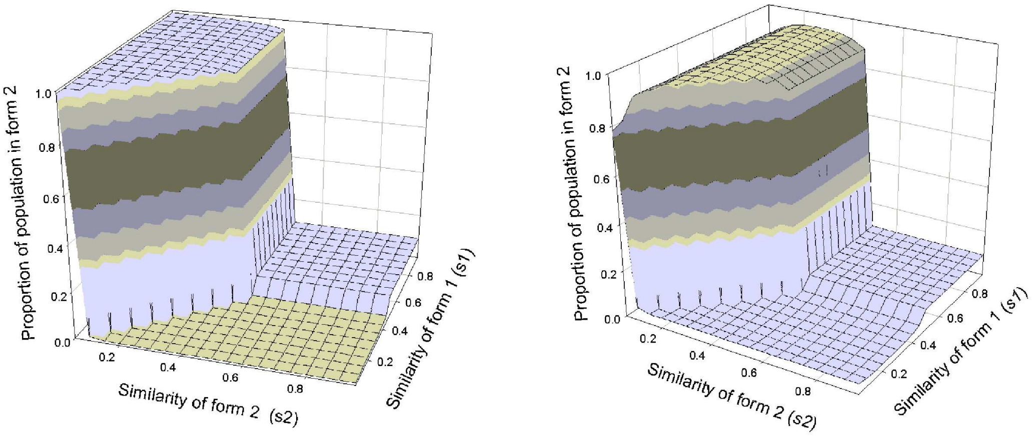

(c) Defended Species

(d) Undefended Species

Figure 8.2: Proportions of defended and undefended prey that evolve to form 2 after $t = 1000$ for a variety of values ($0.05 - 0.95$ in steps of 0.05) of $s1$ and $s2$ *(a-d)*. In *a-b*: $q1 = 0.1$, $q2 = 0.5$ (form 2 is more conspicuous). In *c-d*: $q1 = 0.1$, $q2 = 0.11$ (form 2 is marginally more conspicuous). In all cases equilibria were reached. Parameter values: $b = 2$, $c = 1$, $m = 0.001$, $r = 1$, $K = 1000$, $D1_0 = 500$, $D2_0 = 0$, $U1_0 = 500$, $U2_0 = 0$.

prey from the moderately beneficial prey. Thus, predation is deterred and both species evolve a 50% combination of the two forms, maintained simply by mutational balance (no graph shown).

2. The Parallel Forms Are Entirely Distinguishable

In this second extreme case, the parallel forms are completely distinct $S1 \rightarrow 0$ and $S2 \rightarrow 0$. As in the previous case, we assume that one prey form is slightly more conspicuous than the other. In this case the undefended species always evolves to adopt its most cryptic form regardless of the values of b and c. Since defended prey are distinct from undefended prey, regardless of which form they adopt, optimal predators completely avoid them and thus, they evolve a 50% combination of both forms, maintained by mutational balance.

3. Simultaneously Varying Mimicry and Conspicuousness

In this case we consider the evolution of the two species when one prey form is much more conspicuous than the other ($q1 << q2$), and prey adopting one form are much more easily discriminable than prey adopting the other ($s1 << s2$ or $s1 >> s2$). When the cryptic form (the form with the lowest value of q) is the most discriminable, then both the defended and undefended species evolve to adopt the cryptic form. In contrast, when the conspicuous form is the most discriminable then defended species evolves to adopt the more conspicuous form while the undefended species evolves to adopt the more cryptic form (see Table 8.1).

4. Varying the Degree of Mimicry

We now investigate the implications of the model over a wider range of conditions. Although results here are shown for $b > c$, and when both prey species start initially in form 1, all of the central findings hold under a range of values of $b < c$ and when both species start from form 2 (see Table 8.1).

In this case, both forms of the two prey species could be discriminated to some extent ($0 < s1$, $s2 < 1$) and one form is markedly more conspicuous than the other. Defended prey tended to evolve the more conspicuous form under a broad range of different combinations of $s1$ and $s2$ (Figure 8.1*a*, 8.2*a-b*). On the other hand, undefended species largely remain in their cryptic form (Figure 8.1*b*, 8.2*a-b*). The difference in the form that defended and undefended prey adopt is most prominent when prey species that adopt the conspicuous form are more easily distinguished than when they adopt the cryptic form

($s2 \leq 0.4$ in Figure 8.2*a*-*b*). However, discriminability is not the only factor influencing the outcome. When cryptic prey forms can be more readily distinguished ($s1 < s2$), for example, then defended prey sometimes still evolve to adopt the conspicuous form while undefended prey remain cryptic.

When the more conspicuous form is only marginally more conspicuous than the other form, then the conditions under which defended prey evolve to adopt the conspicuous form is more restricted (Figure 8.2*c*-*d*). Further, under these conditions, when defended prey do evolve to adopt the conspicuous form (under equilibrum), then undefended prey also typically evolve to adopt the marginally more conspicuous form (mimicking the defended prey).

		$b > c \quad [b = 2, c = 1]$	$b < c \quad [b = 1, c = 2]$
$q1 < q2$ $q1 = 0.2$ $q2 = 0.8]$	$s1 < s2$ $[s1 = 0.2, s2 = 0.8]$	Defended: *cryptic* Undefended: *cryptic*	Defended: *cryptic* Undefended: *cryptic*
	$s1 > s2$ $s1 = 0.8, s2 = 0.2$	Defended: *conspicuous* Undefended: *cryptic*	Defended: *conspicuous* Undefended: *cryptic*

Table 8.1: Evolving a means to avoid predation: the equilibrium morphological form exhibited by the vast majority of profitable and unprofitable prey at $t = 1000$ under various combinations of parameters $q1$, $q2$, $s1$, $s2$, b and c. Exact example parameter values given in square brackets. Other parameter values $r = 1$, $K = 1000$, $m = 0.001$. These conclusions remained the same whether all prey start as form 1 ($D1_0 = 500$, $D2_0 = 0$, $U1_0 = 500$, $U2_0 = 0$) or all prey start as form 2 ($D1_0 = 0$, $D2_0 = 500$, $U1_0 = 0$, $U2_0 = 500$).

5. *Varying the Degree of Conspicuousness*

Regardless of the degree of conspicuousness, when prey of form 1 can be much more readily confused than prey of form 2 ($s1 >> s2$) then the defended species tends to evolve form 2 (the less distinguishable form) regardless of its relative degree of conspicuousness (Figure 8.3*a*). Thus, in the unlikely (but sometimes possible, Wüster *et. al.*, 2004) situation where the cryptic form is more easily distinguishable than conspicuous

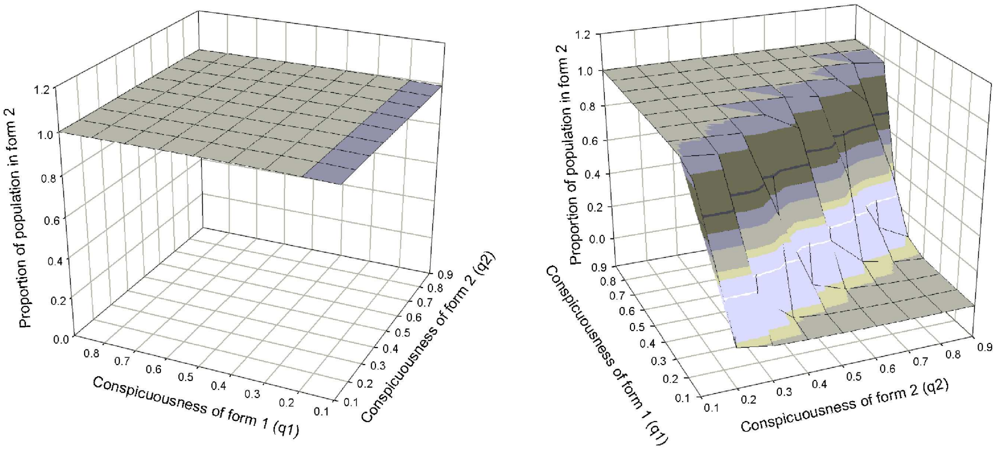

(a) Defended Species

(b) Undefended Species

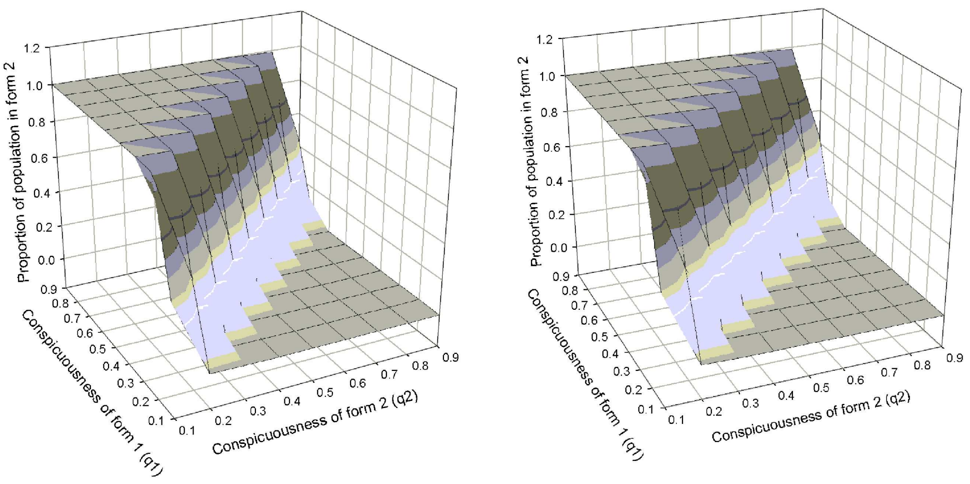

(c) Defended Species

(d) Undefended Species

Figure 8.3: Proportions of defended and undefended prey that evolve to form 2 after $t = 2000$ for a variety of values of $q1$ and $q2$ ($0.1 - 0.9$ in steps of 0.1). In all cases equilibria were reached. Parameter values: In ***a-b***: $s1 = 0.5$, $s2 = 0.1$. In ***c-d***: $s1 = 0.7$, $s2 = 0.7$. In all cases: $b = 2$, $c = 1$, $m = 0.001$, $r = 1$, $K = 1000$, $D1_0 = 500$, $D2_0 = 0$, $U1_0 = 500$, $U2_0 = 0$.

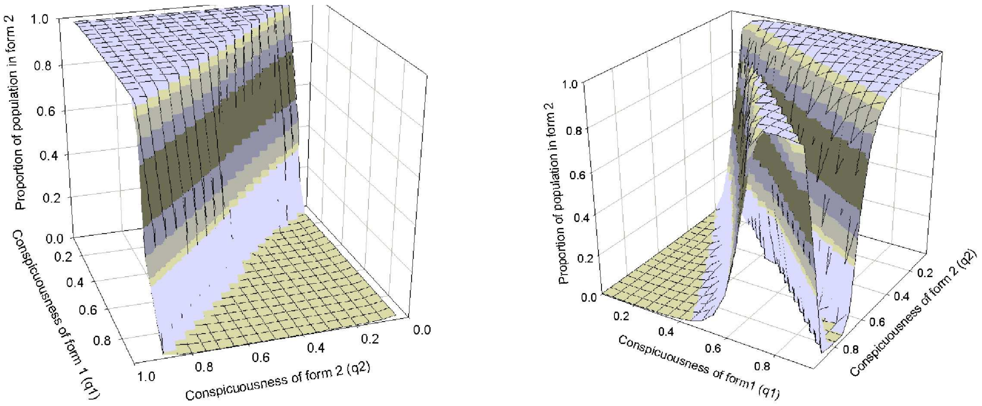

Figure 8.4: Proportions of defended (left) and undefended (right) prey that evolve to form 2 after $t = 1000$ for a variety of values ($0.05 - 0.95$ in steps of 0.05) of $q1$ and $q2$ (the level of conspicuousness). In all cases equilibria were reached. Parameter values: $s1 = 0.1$, $s2 = 0.1$ (both forms are equally discriminable), $b = 2$, $c = 1$, $m = 0.001$, $r = 1$, $K = 1000$, $D1_0 = 500$, $D2_0 = 0$, $U1_0 = 500$, $U2_0 = 0$.

form, defended prey will evolve the more cryptic form. Undefended prey are highly sensitive to conspicuousness, as might be expected, and evolve to adopt the most cryptic form (Figure 8.3*b*).

When parallel forms of the two species are both relatively similar ($s1 = s2 = 0.7$, for example), then both defended and undefended prey evolve to adopt the most cryptic form (Figure 8.3*b*-*c*). However, when the parallel forms are relatively distinct ($s1 = s2 = 0.1$, for example) then defended prey evolve to adopt the most conspicuous form while undefended prey evolve to adopt the most cryptic form. In this case the undefended prey rarely become mimetic (and thus conspicuous) as, in the absence of good mimicry, predators can detect them too easily when conspicuous (Figure 8.4). There are, however, some conditions where the benefits of mimicking the conspicuous display of the defended species outweigh the costs of increased detectability, particularly when the two forms are similar in their level of conspicuousness (Figure 8.4).

6. Model Extension: Three Forms of Each Prey Species

When the model was extended to include three forms of each species, then qualitatively identical results hold. For example, when $b = 2$, $c = 1$, $s1 = s2 = s3 = 0.2$ and $q1 = 0.1$, $q2 = 0.2$, $q3 = 0.5$ then defended species evolve to become predominantly conspicuous while undefended prey evolves to become predominantly cryptic.

8.6 Discussion

Warning signals and mimicry have been intensively studied, and are among the most well-known forms of anti-predator adaptations (Ruxton et al., 2004a). Despite the clear relationship between conspicuous warning signals and mimicry, strikingly little has been done to look at how they interact.

Previous theories and models of Batesian mimicry have all proposed that (with the exception of imperfect mimicry theories) any mutant form that brings the mimic closer in appearance to the model will be favoured by natural selection. Thus, we are left with a continuous coevolutionary race between model and mimic. However, we have shown that when elements of the signal are related to conspicuousness this is not always the case. Our approach does reinforce the explanation for why we should expect Batesian mimics to be close in resemblance to their noxious models. If a defended model species is highly conspicuous then there may be far less scope for the gradual improvement of mimicry after a chance mutation has provided the undefended species with imperfect mimicry, as is typically portrayed.

Our results show that defended prey can indeed be expected to adopt traits in order to shake off undefended Batesian mimics. However, if the quality of the Batesian mimicry is high then it becomes adaptive for the undefended prey to mimic the defended prey regardless of their conspicuousness. As a result, when Batesian mimicry is perfect, the mimetic burden drags both prey towards a cryptic state in order to maximise their protection from predators.

When defended and undefended prey are not completely confused, then defended prey typically succeed in shaking off their mimics by adopting a conspicuous form, while undefended prey must remain cryptic. In this situation conspicuousness is a particularly difficult trait for undefended prey to adopt because, in the absence of perfect mimicry, the cost of increased detectability that comes along with being conspicuousness outweighs the protection gained by the mimicry.

At the start of the chapter, possible opportunity benefits of conspicuousness were mentioned (such as sexual selection or thermoregulation). Note that the model does not include any additional advantages to conspicuousness. Thus, defended prey were not evolving a conspicuous display simply to take advantage of opportunity benefits, but were using it as a way to escape their mimics and improve the reliability of their warning signal. These results are consistent with those of Sherratt and Beatty (2003) who found, using an artificial experimental system where humans were used to predate upon artificial prey, the defended prey evolved colourations that could not be mimicked.

A divergence in appearance between defended and undefended prey (i.e., defended prey evolved conspicuousness while undefended prey evolved crypsis) was most marked when $s1 >> s2$ and $q1 << q2$ or *vice versa*. This finding might be applicable to a variety of natural systems. As Guilford and Dawkins (1991, p. 3) proposed "Aspects of a signal that lead to its being increasingly detectable may also lead to its being increasingly discriminable". $s1 >> s2$ means, roughly speaking, that there are more ways of being conspicuous than there are cryptic. This assumption seems realistic: to be cryptic a prey item has to share colours with the back ground, which are likely to be limited to colours such as greens and browns (for example), whereas it is reasonable to assume that conspicuous prey are not as limited in their possible colourations. Although this assumption allows the model to result in conspicuous defended species and cryptic undefended species under a broader range of conditions—if we drop the assumption, the model still shows this result under a (slightly narrower) range of conditions.

Returning to the question posed earlier: do defended prey evolve traits that undefended prey find difficult to exploit? This model suggests that the answer is yes. Conspicuous traits are difficult for undefended prey to exploit, and thus it can act to reduce the parasitic mimetic burden on defended prey by making undefended prey more vulnerable to predation. Therefore, in the absence of perfect mimicry, conspicuousness acts to 'punish' undefended prey (dishonest exploiters) more than defended prey (honest signallers). As such, the theory shares *similar* characteristics with the theory of minimal-cost, and handicap signalling (although see arguments in Chapter 4).

Although our results support Wallace's theory that defended prey have evolved to distinguish themselves from undefended prey, they do more than just reiterate his argument. For example, the results shown in Figure 8.1*b*, where $s1 = s2$ (both forms are equally discriminable) cannot be explained on the basis of discriminability alone. Further, our analysis suggests that defended prey can, under some conditions, evolve conspicuousness to shake off their mimics even if that form is less distinguishable than an alternative, more cryptic form. However, in (what would probably be rare) cases where the cryptic form can be discriminated more than the conspicuous form, defended prey can sometimes evolve crypsis. This might possibly offer an explanation for the single observation of cryptic defended European vipers which appear to have differentiated themselves from undefended snakes with the invention of zig-zag markings on their back (Wüster, 2004) [2].

Although the discussion in this chapter has, at times, appeared critical of the receiver psychology perspective, it is not intended that the two theories are seen as mutually exclusive alternatives, as both theories are entirely plausible. However, the receiver psychology

---

[2] Another explanation is that the zig-zag pattern appears conspicuous to the particular predators involved.

perspective has in the past been continually championed without having its assumptions questioned, and it is useful to critically examine all theories concerning the evolution of warning signals. Indeed, in nature both of the mechanisms we have considered might have operated in concert. Alternatively, some defended prey might have evolved conspicuous warning signals for the reasons demonstrated in the model, resulting in coevolved predator biases that might then leave the door open for more species to evolve warning signals. Thus, the two theories could possibly make good bedfellows.

In a seminal article that has often been quoted in this thesis, Turner (1984, p. 142) comments on Wallace's argument: "...it may be that aposematic colouring has evolved to take warning patterns away, beyond any possibility of confusion from the green and brown colours of the palatable cryptic prey for which predators are constantly forming search images." With this model, we have offered a mechanism by which this can happen, and helped to elucidate the conditions under which it might be expected. When speaking of how insects might develop primary defences, Turner (1984, p. 1) commented that "...the result is that some insects are not nearly so nice, or so safe, to eat as others. The stage is set for the evolution, first, of warning colouration, and then of mimicry." However, the results here suggest an additional possibility: a stage set, first, for the evolution of mimicry, and then of warning signals.

In summary, warning signals and mimicry are clearly intertwined phenomena. Although isolated models of warning signals and mimicry have set the necessary groundwork for theoretical studies, they might be missing something important. The model presented in this chapter, inspired by the ideas of Wallace, has shown that defended prey may well evolve conspicuous traits to shake off their Batesian mimics. Warning signals can evolve in response to mimicry.

8.7 Future Extensions

The model presented in this chapter is the first to allow for both mimicry and conspicuousness. It is also the first formulation of the idea that conspicuous warning signals have evolved in defended prey to distinguish them from undefended prey. Thus, modelling the theory formally was appropriate. However, as discussed in Chapter 3, mathematical models often benefit from support from evolutionary simulation models that relax assumptions and add extensions. Thus, the intention is to extend this model using an evolutionary simulation model. There are a number of possible extensions that would give us more insight as to the implications of the theory:

- *The coevolutionary-arms-race with conspicuousness as a factor:* Although the current model allows for mimicry and conspicuousness, only the level of conspicuousness is allowed to evolve, whereas the similarity coefficients are fixed. A useful extension would allow a trait independent of conspicuousness to evolve. Although predators may well generalise to a greater extent over cryptic prey than conspicuous prey, the ability to discriminate between prey would thus also depend on the level of evolved mimetic refinement. Thus, in this model prey that are conspicuous, for example, could be readily distinctive or be generalised over depending on their phenotypic distance from another prey species in terms of a dimension unrelated to conspicuousness (as in the mimicry ring model). Thus, a mutation could place an undefended prey item anywhere among a range of conspicuous colourations, and to become a mimic an undefended prey item would need to 'land' by chance on a colouration that is within a close generalization gradient of a conspicuous defended species. As a result, it might be much easier for defended prey to shake off potential mimics. In this situation, what is stopping a defended species evolving back towards a more distinct conspicuous colouration? The next point might suggest an answer.

- *Multi-species model:* The current model only allows for two separate species. An evolutionary simulation modelling approach would allow us to easily include multiple interacting species. There are many questions that such an extension would allow us to address. How does Müllerian mimicry effect the results? What if there are many Batesian mimics? Considering that there are more ways to be conspicuous than cryptic, might undefended prey typically clog the phenotypic landscape of cryptic colourations, making it difficult for defended prey to sustain crypsis?

- *A continuum of conspicuousness:* The current model allows a maximum of three different forms and thus, a maximum of three different levels of conspicuousness. This could be extended to allow the level of conspicuousness to be continuous (thus allowing for the possibility of gradual evolution towards conspicuousness).

  Summers and Clough (2001) provided the first evidence that highly toxic prey tend to be more conspicuous than mildly toxic prey. Thus, allowing a continuum of conspicuousness would allow us to elucidate conditions under which prey with higher levels of toxins might be able to afford to evolve more highly conspicuous traits than prey with less fierce toxins.

- *Other traits that might be difficult to exploit:* The current model examines the implications of conspicuous colourations. This could be extended to look at other traits such as aggregations, slow movement, and so forth, and evaluate the robustness of the argument for a far wider set of circumstances.

- *Multi-modal traits:* Joron (2003) suggests that multi-modal signals (i.e., a display that is simultaneously conspicuous in multiple sensory modalities) have evolved in defended species in an attempt to escape mimicry from undefended species. Multiple signalling dimensions, in which more than one dimension is correlated with conspicuousness, could be implemented in the model to examine this idea.

Chapter 9

Conclusions

9.1 Biological Implications

The following conclusions can be taken from this thesis:

- Evolutionary simulation models, in support of mathematical models, are a useful tool for modelling the evolution of warning signals and mimicry. This type of model has much to offer to theoretical studies and take them forwards.
- A model presented in Sherratt (2002a), although inspiring and novel, does not predict the evolution of warning signals. The model allows for two fixed populations: one conspicuous and one cryptic, each of which can evolve varying levels of defence. However, the typical approach is to model two populations where one is defended and the other undefended. Thus, the evolving trait is their level of conspicuousness/crypsis. Sherratt's model always results in conspicuous defended prey, undefended cryptic prey, and predators that disproportionally attack cryptic prey. Thus, it is easily misinterpreted as explaining the evolution of warning signals—but this is not the case. If the level of conspicuousness is allowed to evolve along with the level of defence then, in the model, both defended and undefended prey evolve to become cryptic. Thus, the work here has clarified exactly what it is that the model implies.
- Diverse mimicry rings can form because intermediate mutational forms (mutant

prey that do not share the appearance of any ring) are at a selective disadvantage. This theory is more general and parsimonious than competing (or additional) theories that prey are being constrained by sexual selection, or that mimicry rings are diverse due to minor variations in prey flight or roosting height.

- Batesian mimics influence the evolution of Müllerian mimicry rings. The presence of Batesian mimics can effect the evolution of mimicry rings. Batesian mimics can parasitise rings, causing them to destabilise. This can stop rings from stabilizing and thus, cause them to converge with neighbouring rings. This convergence of mimicry rings can subsequently cause them to stabilise, and there are enough defended species in the ring to be able to tolerate the Batesian mimic. Thus, Batesian mimics can cause a reduction in the number of mimicry rings in an ecosystem.

- Warning signals can evolve in response to mimicry. If both warning signals and mimicry are modelled together, then defended prey can shake-off their parasitic Batesian mimics by adopting traits that undefended prey find difficult to exploit—such as conspicuous warning signals. Thus, warning signals may well have evolved for reasons similar to those originally given by Wallace: that defended prey have evolved to distinguish themselves from undefended prey. This contrasts with the conventional explanation.

- Typically, the evolution of warning signals is thought to set the stage for the evolution of mimicry. However, mimicry can occur in species that are not conspicuous, and it has been shown here that mimicry might instead set the stage for the evolution of warning signals.

- Throughout this thesis, two apparently competing theories have been discussed regarding the evolution of warning signals: what I have called the receiver psychology approach and the reliable indicator of defence approach. These approaches have both been examined with a crytical eye. The receiver psychology approach has been impunged disproportionately, but this is probably due to the dogma surrounding this well-developed theory. Indeed, both theories are entirely plausible and may work well in consort.

9.2 Future Work

Naturally, a Ph.D. cannot be an entirely exhaustive work that completely covers an area. Thus, although the work presented here has had a significant impact, this thesis is intended to kick-start the use of evolutionary simulation models for exploring theories of warning signals and mimicry. There are too many options for future work to list them all, but important possibilities include:

- The model of coevolving predators and prey presented in Chapter 4 is loaded with potential. One of the most interesting extensions to this model would allow for the evolution of predator educational biases.

- A related model in which receiver biases could evolve would be useful to explicitly explore the argument that observations of biases in today's predators are not necessarily evidence that warning signals initially evolved to exploit such biases. In a virtual world that contains defended conspicuous prey and undefended cryptic prey, it would be possible to explore the corresponding evolution of predator educational biases.

- Given that Batesian mimicry occurs so readily in the mimicry ring model, why is it not more common in nature? The mimicry ring model presented in Chapter 7 does not allow for evolution along a dimension related to conspicuousness/crypsis. This is not a problem for the model, as we have chosen to examine the effect of Batesian mimics on mimicry rings and not the tradeoff between mimicry and crypsis. Allowing for crypsis in a future model would allow the implications of the model to be assessed in more detail. However, theories of the evolution of warning signals in the presence of mimicry need to be assessed in more detail before this can be reliably done. This brings us on to the next point.

- The mathematical model of warning signals in response to mimicry needs to be extended using an evolutionary simulation model. There are many possible extensions listed in Chapter 8; however, the most important extensions would be to allow the accuracy of the mimicry to evolve and to allow for a richer multi-species ecosystem.

- The previous chapter presented the first attempt at combining theories of warning signals and mimicry. A long-term objective is to include another highly related phenomena into the model—the evolution of primary defences. Such a model would give us a full story of how conspicuous defended and cryptic undefended prey came about.

- Although many models of mimicry allow for learning, many models (including those examining the dynamics) do not. It would be useful to see if differences in the relevant time frames between different predators caused by allowing life-time learning (or lack thereof), or by allowing longer predator generations, and so forth, make a significant difference to mimicry dynamics.

- A system for visualising the evolution of morphologies of artificial prey might be useful for improving our intuitions and understanding of the associated coevolutionary dynamics. Such a system would allow the user to adjust parameters on-the-fly and graphically observe the results.

9.3 Broader Implications

It is worth understanding warning signals and mimicry for their own sakes, as they are complex natural phenomena that have intrigued people for a long time. However, the work does have broader implications. Developing a better understanding of coevolutionary phenomena in general, such as how to model them, what classes of dynamics to expect, and so forth, is useful to evolutionary biology in general. Mimicry, with its interesting coevolutionary dynamics, and warning signals with their associated predator-prey interactions are useful sources of information for coevolution.

Furthermore, a full understanding of coevolution means more than just understanding the natural world, but also will mean almost limitless applications in evolutionary computing (see, e.g., Cartlidge 2004). Clearly this thesis did not focus on the use of coevolution as an optimization tool. However, the major research problems in coevolutionary optimization involve attempts to maintain "productive" dynamics between co-evolving populations. Watson and Pollack (2001) outline problems such as disengagement, overspecialization, and cycling: all cases in which the coevolving populations "get stuck" in one way or another (for more detail see Cartlidge, 2004). Solving these problems will involve new ideas in visualizing and influencing coevolutionary progress in a highly multidimensional space of strategies. This thesis has been concerned with similar problems when tracking and explaining the coevolutionary trajectories of Batesian and Müllerian mimics, and thus should be a rich source of material for the coevolutionary engineer.

Warning signals and mimicry may also have implications for signalling theory. The two phenomena can fit into signalling categories, such as handicap signalling and conventional signalling, to a limited extent. However, the fit is awkward. Mimicry and warning signals are a good example of why we need to expand our theoretical vocabulary when talking about the ways in which one animal influences the behaviour of another (Noble, 1998).

In conclusion, warning signals and mimicry are interesting evolutionary phenomena which can shed light on many interesting problems in (co)evolution. This thesis has shown that evolutionary simulation modelling, in addition to traditional techniques, is a useful theoretical tool for exploring these well-known exemplars of coevolution.

Bibliography

Adams, E. S., & Caldwell, R. L. (1990). Deceptive communication in asymmetric fights of the stomatopod crustacean *Gonodactylus bredini*. *Animal Behaviour*, *39*, 706–716.

Alatalo, R. V., & Mappes, J. (1996). Tracking the evolution of warning signals. *Nature*, *382*, 708–710.

Azmeh, S., Owen, J., Sorensen, K., Grewcock, D., & Gilbert, F. (1998). Mimicry profiles are affected by human-induced habitat changes. *Proceedings of the Royal Society of London: Biological Sciences*, *265*, 2285–2290.

Bates, H. W. (1862). Contributions to an insect fauna of the Amazon valley. *Lepidoptera: Heliconidae*. *Transactions of the Linnean Society*, *23*, 495–566.

Bates, H. W. (1863). *The Naturalist on the River Amazons*. John Murray, London.

Beatty, C. D., Beirinckx, K., & Sherratt, T. N. (2004). The evolution of Müllerian mimicry in multispecies communities. *Nature*, *431*, 63–66.

Beccaloni, G. W. (1997). Vertical stratification of ithomiine butterfly *(Nymphalidae: Ithomiinae)* mimicry complexes: The relationship between adult flight height and larval host-plant height. *Biological Journal of the Linnean Society*, *62*, 313–341.

Bhaskar, R. (1978). *A realist theory of science*. Harvester Wheatshead, Hemel Hempstead, UK.

Blakiston, T., & Alexander, T. (1884). Protection by mimicry—a problem in mathematical zoology. *Nature*, *29*, 405–406.

Bobisud, L. E., & Potratz, J. C. (1976). One-trail versus multi-trial learning for a predator encountering a model-mimic system. *American Naturalist*, *110*, 121–128.

Boden, E. M. A. (1996). *The Philosophy of Artificial Life*. Oxford University Press, Oxford.

Bonabeau, M. A. (1994). Why do we need artificial life? *Artificial Life*, *1*(2), 303–325.

Bowler, P. J. (1944). *Evolution: The History of an Idea*, Vol. 3. University of California Press.

Brower, J. V. (1958a). Experimental studies of mimicry in some North-American butterflies .1. the monarch, *Danaus plexippus*, and viceroy, *Limenitis archippus-Archippus*. *Evolution*, *12*, 32–47.

Brower, J. V. (1958b). Experimental studies of mimicry in some North-American butterflies .2. *Battus philenor* and *Papilio troilus, P. polyxenes* and *P. glaucus*. *Evolution*, *12*, 123–136.

Brower, J. V. (1958c). Experimental studies of mimicry in some North-American butterflies .3. *Danaus gilippus berenice* and *Limenitis archippus floridensis*. *Evolution*, *12*, 273–285.

Brower, J. V. (1960). Experimental studies of mimicry .IV. The reactions of starlings to different proportions of models and mimics. *American Naturalist*, *94*, 271–282.

Brower, L. P., Alcock, J., & Brower, J. V. Z. (1971). Avian feeding behavior and the selective advantage of incipient mimicry. In Creed, R. (Ed.), *Ecological Genetics and Evolution, Essays in Honour of E.B. Ford*, pp. 261–274. Blackwell Scientific Publications, Oxford, England.

Brower, L. P., & Brower, J. V. Z. (1972). Parallelism, convergence, divergence, and the new concept of advergence in the evolution of mimicry. In Deevey, E. A. (Ed.), *Ecological essays in honour of G. Evelyn Hutchinson*, Vol. 44, pp. 59–67. Transactions of the Connecticut Academy of Science.

Brower, L. P., Ryerson, W. N., Coppinger, W. N., & Glazier, S. (1968). Ecological chemistry and the palatability spectrum. *Science*, *161*, 1349–1351.

Bullock, S. (1997). *Evolutionary simulation models: On their character and application to problems concerning the evolution of natural signalling systems*. Ph.D. thesis, School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK.

Bullock, S. (1999). Are artificial mutation biases unnatural? In Floreano, D., Nicoud, J.-D., & Mondada, F. (Eds.), *Fifth European Conference on Artificial Life (ECAL99)*, pp. 64–73. Springer, Heidelberg.

Burd, M. (1994). Butterfly wing colour patterns and flying heights in the seasonally wet forest of barro colorado island, panama. *Journal of Tropical Ecology*, *10*, 601–610.

Cartlidge, J. (2004). *Rules of Engagement: Competitive coevolutionary dynamics in computational systems*. Ph.D. thesis, School of Computing, Univeristy of Leeds.

Chai, P. (1986). Field observations and feeding experiments on the responses of rufous-tailed Jacamars *(Galbula rufi cauda)* to free-flying butterflies in a tropical rain-forest. *Biological Journal of the Linnean Society*, *29*, 161–189.

Chalmers, A. F. (1999). *What Is This Thing Called Science?* (3rd edition). Open University Press, Buckingham.

Charlesworth, D., & Charlesworth, B. (1975). Theoretical genetics of Batesian mimicry. *Journal of Theoretical Biology*, *55*, 283–303.

Cheney, D. L., & Seyfarth, R. M. (1982). How vervet monkeys perceive their grunts: Field playback experiments. *Animal Behaviour*, *30*, 739–751.

Cheney, D. L., & Seyfarth, R. M. (1990). *How Monkeys See the World*. Chicago University Press, Chicago.

Clarke, C. A., & Sheppard, P. M. (1971). Further studies on the genetics of the mimetic butterfly *Papilio memmon*. *Philosophical Transactions of the Royal Society*, *263*, 431–458.

Cook, L. M. (2000). Changing views on melanic moths. *Biological Journal of the Linnean Society*, *69*, 431–441.

Coppinger, R. (1969). The effect of experience and novelty on avian feeding behaviour with reference to the evolution of warning coloration in butterflies. Part I: reactions of wild-caught adult blue jays to novel insects. *Behaviour*, *35*, 45–60.

Coppinger, R. (1970). The effect of experience and novelty on avian feeding behaviour with reference to the evolution of warning coloration in butterflies. Part II: reactions of naïve birds to novel insects. *American Naturalist*, *104*, 323–335.

Cott, H. B. (1940). *Adaptive Colouration in Animals*. Methuen, London.

Darwin, C. (1859). *The Origin of Species by Means of Natural Selection or the Preservation of Favoured Races in the Struggle for Life*. John Murray, London.

Darwin, C. (1871). *The Descent of Man and Selection in Relation to Sex*. John Murray, London.

Darwin, C. (1877). *The Various Contrivances by which Orchids are Fertilised by Insects* (2nd edition). Murray, London.

Darwin, C. (1887). *The Life and Letters of Charles Darwin: Including an Autobiographical Chapter, Edited by his Son Francis Darwin*. Murray, London.

Dawkins, R., & Krebs, J. R. (1979). Arms races between and within species. *Proceedings of the Royal Society of London B: Biological Sciences*, *205*, 489–511.

DeVries, P. J., Lande, R., & Murray, D. (1999). Associations of co-mimetic ithomiine butterflies on small spatial and temporal scales in a neotropical rainforest. *Biological Journal of the Linnean Society*, *67*, 73–85.

Di Paolo, E. A. (1999). *On the Evolutionary and Behavioral Dynamics of Social Coordination: Models and Theoretical Aspects*. Ph.D. thesis, School of cognitive and computing sciences, University of Sussex.

Di Paolo, E., Noble, J., & Bullock, S. (2000). Simulation models as opaque thought experiments. In Bedau, M. A., McCaskill, J. S., Packard, N. H., & Rausmussen, S. (Eds.), *Artificial Life VII: The Seventh International Conference on the Simulation and Synthesis of Living Systems*, pp. 497–506. MIT Press/Bradford Books, Cambridge, MA.

Dixey, F. A. (1909). On Müllerian mimicry and diaposematism. *Entomological Society of London*, *XXII*, 559–583.

Douglas, M. M. (1979). Hot butterflies. *Natural History*, *88*(9), 56–65.

Douwes, P. (1976). Activity in *Heodes virgaureae* in relation to air temperature, solar radiation, and time of day. *Oecologia*, *22*, 287–298.

Dukas, R., & Ellner, S. (1993). Information processing and prey detection. *Ecology*, *74*, 1337–1346.

Edmunds, M. (2000). Why are there good and poor mimics? *Biological Journal of the Linnean Society*, *13*, 605–618.

Ehrlich, P. R., & Raven, P. H. (1964). Butterflies and plants: A study in co-evolution. *Evolution*, *18*, 586–608.

Emlen, J. M. (1968a). Batesian mimicry—a preliminary theoretical investigation od quantitative aspects. *American Naturalist*, *102*, 235–241.

Emlen, J. M. (1968b). Batesian mimicry—a preliminary theoretical investigation of quantitative aspects. *American Naturalist*, *102*, 235–241.

Endler, J. A. (1978). A predator's view of animal color patterns. *Evolutionary Biology*, *11*, 25–31.

Endler, J. A. (1988). Frequency-dependent predation, crypsis and aposematic coloration. *Philosophical Transactions of the Royal Society of London: Biological Sciences*, *319*, 505–523.

Endler, J. A. (1991). Interactions between predators and prey. In Krebs, J. R., & Davies, N. B. (Eds.), *Behavioural Ecology. 3rd ed.*, pp. 169–196. Blackwell Scientific, Oxford.

Endler, J. A., & Mappes, J. (2004). Predator mixes and the conspicuousness of aposematic signals. *The American Naturalist*, *163*(4), 532–547.

Enquist, M. (1985). Communications during aggressive interactions with particular reference to variation in choice of behaviour. *Animal Behaviour*, *33*, 1152–1161.

Estabrook, F. G., & Jespersen, D. C. (1974). Strategy for a predator encountering a model-mimic system. *The American Naturalist*, *108*, 443–457.

Feltmate, B. W., & Williams, D. D. (1989). A test of crypsis and predation avoidance in the stonefly paragnetina media (plecoptera: Perlidae). *Animal Behaviour*, *37*, 992–999.

Fisher, R. A. (1930). *The Genetical Theory of Natural Selection*. Oxford: Clarendon Press.

Forbes, J. M. (2001). Consequences of feeding for future feeding. *Comparative Biochemistry and Physiology*, *128*(Part A), 463–470.

Gamberale, G., & Tullberg, B. S. (1996). Evidence for a peak-shift in predator generalization among aposematic prey. *Proceedings of the Royal Society of London: Biological Sciences*, *263*, 1329–1334.

Gavrilets, S., & Hastings, A. (1998). Coevolutionary chase in two-species systems with applications to mimicry. *Journal of Theoretical Biology*, *191*, 415–427.

Gendrom, R. P., & Staddon, J. E. R. (1983). Searching for cryptic prey: the effects of search rate. *American Naturalist*, *121*, 172–186.

Getty, T. (1985). Discriminability and the sigmoid functional response — how optimal foragers could stabilize model-mimic complexes. *American Naturalist*, *125*, 239–256.

Getty, T., Kamil, A. C., & Real, P. G. (1987). Signal detection theory and foraging for cryptic or mimetic prey. In Kamil, A., Krebs, J. R., & Pulliam, H. R. (Eds.), *Foraging Behavior*, pp. 525–549. Plenum Press, New York.

Gillette, E., Huang, R.-C., Hatcher, N., & Moroz, L. (2000). Cost-benefit analysis potential in feeding behavior of a predator snail by integration of hunger, taste, and pain. *Proceedings of the National Academy of Science USA*, *97*, 3585–3590.

Gittleman, J. L., & Harvey, P. H. (1980). Why are distasteful prey not cryptic? *Nature*, *28*, 897–899.

Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley.

Goldschmidt, R. B. (1945). Mimetic polymorphism, a controversial chapter of darwinism. *Quarterly Review of Biology*, *20*, 147–164.

Grafen, A. (1982). How not to measure inclusive fitness. *Proceedings of the Royal Society of London: Biological Sciences*, *298*, 425.

Grafen, A. (1990). Biological signals as handicaps. *Journal of Theoretical Biology*, *144*, 517–546.

Greenwood, J. J. D. (1986). Crypsis, mimicry, and switching by optimal foragers. *American Naturalist*, *128*, 294–300.

Guilford, T. (1986). How do ‘warning colors’ work? Conspicuousness may reduce recognition errors in experienced predators. *Animal Behaviour*, *34*, 286–288.

Guilford, T. (1988). The evolution of conspicuous coloration. *American Naturalist*, *131*, 7–21.

Guilford, T. (1990). The evolution of aposematism. In Evans, D. L., & Schmidt, J. O. (Eds.), *Insect Defenses: Adaptive Mechanisms and and Strategies of Prey and Predators*, pp. 23–61. University of New York Press.

Guilford, T. (1994). ‘Go-slow’ signalling and the problem of automimicry. *Journal of Theoretical Biology*, *170*, 311–316.

Guilford, T., & Dawkins, M. S. (1987). Search images not proven: A reappraisal of recent evidence. *Animal Behaviour*, *35*, 1838–1845.

Guilford, T., & Dawkins, M. S. (1991). Receiver psychology and the evolution of animals signals. *Animal Behaviour*, *42*, 1–14.

Guilford, T., & Dawkins, M. (1993). Are warning colors handicaps? *Evolution*, *47*, 400–416.

Hadeler, K. P., Demottoni, P., & Tesei, A. (1982). Mimetic gain in Batesian and Müllerian mimicry. *Oecologia*, *53*, 84–92.

Hammerstein, P. (1996). Darwinian adaptation, population genetics and the streetcar theory of evolution. *Journal of mathematical biology*, *34*(5-6), 511–532.

Hanson, H. M. (1959). Effects of discrimination training on stimulus generalization. *Journal of Experimental Psychology*, *58*, 321–334.

Härlin, C., & Härlin, M. (2003). Towards a historization of aposematism. *Evolutionary Ecology*, *17*, 197–212.

Harvey, P. H., Bull, J. J., Pemberton, M., & Paxton, R. J. (1982). The evolution of aposematic coloration in distasteful prey: a family model. *American Naturalist*, *119*, 710–719.

Harvey, P. H., & Paxton, R. J. (1981). The evolution of aposematic coloration. *Oikos*, *27*, 391–396.

Hasson, O. (1994). Cheating signals. *Journal of Theoretical Biology*, *167*, 223–238.

Heiling, A. M., Herberstein, M. E., & Chittka, L. (2003). Crab-spiders manipulate flower signals. *Nature*, *421*, 331.

Hileman, K. S., Brodie, E. D., & Formanowicz, D. R. (1995). Avoidance of unpalatable prey by predaceous diving beetle larvae—the role of hunger level and experience *(Coleoptera, Dytiscidae)*. *Journal of Insect Behaviour*, *8*, 241–249.

Hillis, W. D. (1990). Co-evolving parasites improve simulated evolution as an optimization procedure. *Physica D*, *42*, 228–234.

Holling, C. S. (1965). The functional response of predators to prey density, and its role in mimicry and population regulation. *Memoirs of the Entomological Society of Canada*, *45*, 1–60.

Holmgren, N., & Enquist, M. (1999). Dynamics of mimicry evolution. *Biological Journal of the Linnean Society*, *66*, 145–158.

Huheey, J. E. (1964). Studies of warning coloration and mimicry. IV. A mathematical model of model-mimic frequencies. *Ecology*, *45*, 185–188.

Huheey, J. E. (1976). Studies in warning coloration and mimicry. VII. Further evidence for a frequency-dependent model of predation. *Journal of Herpetology*, *14*, 223–230.

Hurd, P. L., & Enquist, M. (1998). Conventional signalling in aggressive interactions: the importance of temporal structure. *Journal of Theoretical Biology*, *192*, 197–211.

Huxley, J. (1942). *Evolution, the Modern Synthesis*. Allen and Unwin, London.

Huynen, M., Stadler, P., & Fontana, W. (1996). Smoothness within ruggedness: The role of neutrality in adaptation. *Proceedings of the National Academy of Science USA*, *93*(1), 397–401.

Janzen, D. H. (1980). When is it co-evolution? *Evolution*, *34*(3).

Jiggins, C. D., McMillan, W. O., Neukirchen, W., & Mallet, J. (1996). What can hybrid zones tell us about speciation? The case of Heliconius erato and H. himera *(Lepidoptera: Nymphalidae)*. *Biological Journal of the Linnean Society*, *59*, 221–242.

Johannesson, K., & Ekendahl, A. (2002). Selective predation favouring cryptic individuals of marine snails (*Littorina*). *Biological Journal of the Linnean Society*, *76*(137-144).

Johnstone, R. (2002). The evolution of innacurate mimics. *Nature*, *418*, 524–526.

Joron, M. (2003). Aposematic coloration. In Cardé, R., & Resh, V. (Eds.), *Encyclopedia of Insects*, pp. 39–45. Academic Press, New York.

Joron, M., Wynee, I. R., Lamas, G., & Mallet, J. (2001). Variable selection and the coexistence of multiple mimetic forms of the butterfly *Heliconius numata*. *Evolutionary Ecology*, *13*, 721–754.

Kannan, D. (1983). A markov-chain analysis of predator strategy in a model-mimic system. *Bulletin of Mathematical Biology*, *45*, 347–400.

Kokko, H., Mappes, J., & Lindstrom, L. (2003). Alternative prey can change model-mimic dynamics between parasitism and mutualism. *Ecology Letters*, *6*, 1068–1076.

Krebs, J. R., & Davies, N. B. (Eds.). (1997). *Behavioural Ecology: An Evolutionary Approach* (Fourth edition). Blackwell Science Ltd.

Krebs, J. R., & Dawkins, R. (1984). Animal signals: Mind reading and manipulation. In Krebs, J. R., & Davies, N. B. (Eds.), *Behavioural Ecology: An Evolutionary Approach* (2nd edition)., pp. 380–402. Blackwell, Oxford.

Lamarck, J. (1809). *Philosophie zoologique, ou exposition des considérations relatives à l'historie naturalle des animaux*. E. Dentu, Paris.

Langton, C. G. (Ed.). (1989). *Artificial Life*. Addison-Wesley, Redwood City, CA.

Leimar, O., Enquist, M., & Sillén-Tullberg, B. (1986). Evolutionary stability of aposematic coloration and prey unprofitability: A theoretical analysis. *American Naturalist*, *128*, 469–490.

Levy, S. (1993). *Artificial Life*. Vintage Books, New York.

Lindström, L. (1999). Experimental approaches to studying the initial evolution of conspicuous aposematic signalling. *Evolutionary Ecology*, *13*, 605–618.

Lindstrom, L., Alatalo, R. V., Lyytinen, A., & Mappes, J. (2001). Strong antiapostatic selection against novel rare aposemastic prey. *Proceedings of the National Academy of Sciences of the USA*, *98*, 9181–9184.

Luedeman, J. K., McMorris, F. R., & Warner, D. D. (1981). Predators encountering a model-mimic system with alternative prey. *American Naturalist*, *117*, 1040–1048.

MacDougall, A., & Dawkins, M. S. (1999). Predator discrimination error and the benefits of Müllerian mimicry. *Animal Behavior*, *55*, 1281–1288.

Mackintosh, N. J. (1985). *The psychology of animal learning*. Academic Press, London.

Malcolm, S. B. (1990). Mimicry: status of a classical evolutionary paradigm. *Trends in Ecology and Evolution*, *5*, 57–62.

Mallet, J. (2004a). Home page.. http://abacus.gene.ucl.ac.uk/jim/Mim/Muller.html.

Mallet, J. (2004b). The peppered moth: a black and white story after all. *Genetics Society News*, *50*, 34–38.

Mallet, J., & Gilbert, L. E. (1995). Why are there so many mimicry rings? Correlations between habitat, behaviour and mimicry in *Heliconius* butterflies. *Biological Journal of the Linnean Society*, *55*, 159–180.

Mallet, J., & Joron, M. (1999). The evolution of diversity in warning color and mimicry: polymorphisms, shifting balance, and speciation. *Annual Review of Ecology and Systematics*, *30*, 201–233.

Mallet, J., & Joron, M. (2000). The evolution of diversity in warning colour and mimicry. *Annual Review of Ecological Systems*, *30*, 201–233.

Mallet, J., & Singer, M. C. (1987). Individual selection, kin selection, and the shifting balance in the evolution of warning signals — the evidence from butterflies. *Biological Journal of the Linnean Society*, *32*, 337–350.

Marshall, G. A. K. (1908). On diaposematism, with reference to some limitations of the Müllerian hypothesis of mimicry. *Transactions of the Entomological Society of London*, *1908*, 93–142.

Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.

Maynard Smith, J., & Harper, D. G. C. (1988). The evolution of aggression: can selection generate viability? *Philosophical Transactions of the Royals Society of London B*, *319*, 557–570.

Maynard Smith, J., & Harper, D. G. C. (2003). *Animal Signals*. Oxford University Press, Oxford.

Mayr, R., & Provine, W. B. (Eds.). (1980). *The Evolutionary Synthesis: Perspectives on the unification of Biology*. Harvard University Press, Cambridge MA.

Miller, G. F. (1995). Artificial life as theoretical biology: How to do real science with computer simulation. Tech. rep., (Cognitive Science Research Paper 378), University of Sussex, Brighton, UK.

Mitchell, M. (1998). *An Introduction to Genetic Algorithms*. The MIT Press.

Mode, C. J. (1958). A mathematical model for the co-evolution of obligate parasites and their hosts. *Evolution*, *12*, 158–165.

Müller, F. (1879). *Ituna* and *Thyridia*: a remarkable case of mimicry in butterflies. *Transactions of the Entomological Society of London*, xx–xxix.

Nicholson, A. J. (1927). Presidential address. a new theory of mimicry in insects. *Australian Zoologist*, *5*, 10–24.

Noble, J. (1998). *The Evolution of Animal Communication Systems: Questions of Function Examined through Simulation*. Ph.D. thesis, School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK.

Noble, J. (2000). Talk is cheap: Evolved strategies for communication and action in asymmetrical animal contests. In Meyer, J.-A., Berthoz, A., Floreano, D., Roitblat, H., & Wilson, S. (Eds.), *From Animals to Animats 6: Proceedings of the Sixth International Conference on the Simulation of Adaptive Behavior*, pp. 481–490 Cambridge, MA. MIT Press.

Nur, U. (1970). Evolutionary rates of models and mimics in Batesian mimicry. *American Naturalist*, *104*, 477–486.

Oaten, A., Pearce, C. E. M., & Smyth, B. E. M. (1975). Batesian mimicry and signal-detection theory. *Bulletin of Mathematical Biology*, *37*, 367–387.

Osborn, H. F. (1929). *From the Greeks to Darwin: The Development of the Evolution Idea through Twenty-Four Centuries*. C. Scribner's and Sons.

Owen, A. R. G., & Owen, R. E. (1984). Mathematical paradigms for mimicry—recurrent sampling. *Journal of Theoretical Biology*, 217–247.

Papageorgis, C. (1975). Mimicry in neotropical butterflies. *American Scientist*, *63*, 522–532.

Pietrewica, A. T., & Kamil, A. C. (1975). Search image formation in the blue jay *(Cyanocitta cristata)*. *Nature*, *204*, 1332–1333.

Plaisted, K. C., & Mackintosh, N. J. M. (1995). Visual search for cryptic stimuli in pigeons: implications for the search image and search rate hypothesis. *Animal Behaviour*, *50*, 1219–1232.

Plowright, R. C., & Owen, R. E. (1980). The evolutionary significance of bumble bee color patterns — a mimetic interpretation. *Evolution*, *34*, 622–636.

Pough, F. H., Brower, L. P., Meck, H. R., & Kessell, S. R. (1973). Theoretical investigations of automimicry: Multiple trial learning and the palatability spectrum. *Proceedings of the National Academy of Sciences, U.S.A.*, *70*, 2261–2265.

Poulton, E. B. (1890). *The Colours of Animals: Their Meaning and Use Especially Considered in the Case of Insects*. London: Keegan, Paul, Trench, Trübner.

Poulton, E. B. (1909). *Charles Darwin and the Origin of Species*. Longmans, Green and Co, London.

Prusinkiewicz, P., & Hanan, J. (1989). *Lindenmayer Systems, Fractals, and Plants*. Springer, New York.

Punnett, R. C. (1915). *Mimicry in Butterflies*. Cambridge University Press, London.

Ridley, M. (1993). *The Red Queen. Sex and the Evolution of Human Nature*. Penguin Books Ltd.

Ritland, D. B. (1991). Revising a classic butterfly mimicry scenario: demonstration of Müllerian mimicry between Florida viceroys *(Danaus gilippus berenice)*. *Evolution*, *45*, 918–934.

Roper, T. J. (1994). Conspicuousness of prey retards reversal of learned avoidance. *Oikos*, *69*, 115–118.

Roper, T. J., & Wistow, R. (1986). Aposematic colouration and avoidance learning in chicks. *Quarterly Journal of Experimental Psychology*, *38B*(141-149).

Rowe, C., & Guilford, T. (2000). Aposematism: To be red or dead. *Trends in Ecology and Evolution*, *15*, 261–262.

Ruxton, G. D. (1998). Mimicry—sheep in wolves' clothing. *Nature*, *394*(6696), 833–834.

Ruxton, G. D., Sherratt, T. N., & Speed, M. P. (2004a). *Avoiding Attack: The Evolutionary Ecology of Crypsis, Warning signals & Mimicry*. Oxford University Press, Oxford.

Ruxton, G. D., Speed, M. P., & Sherratt, T. N. (2004b). Evasive mimicry: when (if ever) could mimicry based on difficulty of capture evolve? *Proceedings of the Royal Society of London: Biological Sciences*, *271*(1553), 2135–2142.

Sasaki, A., Kawaguchi, I., & Yoshimori, A. (2002). Spatial mosaic and interfacial dynamics in a Müllerian mimicry system. *Theoretical Population Biology*, *61*, 49–71.

Searcy, W. A. (1992). Song repertoire and mate choice in birds. *American Zoologist*, *32*, 507–533.

Servedio, M. R. (2000). The effects of predator learning, forgetting, and recognition erros on the evolution of warning coloration. *Evolution*, *54*, 751–763.

Sexton, O. J., Hoger, C., & Ortleb, E. (1966). Anolis carolinensis: effects of feeding on reactions to aposematic prey. *Science*, *153*(1140).

Sheppard, P. M., Turner, J. R. G., Brown, K. S., Benson, W. W., & Singer, M. C. (1985). Genetics and the evolution of Muellerian mimicry in *Heliconius* butterflies. *Philosophical Transactions of the Royal Society of London Series B—Biological Sciences*, *308*, 433–610.

Sherratt, T. N. (2001). The evolution of female-limited polymorphisms in damselflies: a signal detection model. *Ecology Letters*, *4*, 22–29.

Sherratt, T. N. (2002a). The coevolution of warning signals. *Proceedings of the Royal Society of London: Biological Sciences*, *269*, 741–746.

Sherratt, T. N. (2002b). The evolution of imperfect mimicry. *Behavioral Ecology*, *13*, 821–826.

Sherratt, T. N. (2003). State-dependent risk-taking in systems with defended prey. *Oikos*, *103*, 93–100.

Sherratt, T. N., & Beatty, C. D. (2003). The evolution of warning signals as reliable indicators of prey defense. *The American Naturalist*, *162*(4), 377–389.

Sherratt, T. N., Speed, M. P., & Ruxton, G. D. (2004). Natural selection on unpalatable species imposed by state-dependent foraging behaviour. *Journal of Theoretical Biology*, *228*, 217–226.

Sherratt, T., & Franks, D. (2004). Do unprofitable prey evolve traits that profitable prey find difficult to exploit? *Submitted to: Proceedings of the Royal Society of London: Biological Sciences*.

Shettleworth, S. J. (1998). *Cognition, Evolution, and Behavior*. Oxford University Press.

Silk, J. B., Kaldor, E., & Boyd, R. (2000). Cheap talk when interests conflict. *Animal Behaviour, 59*, 423–432.

Sillén-Tullberg, B. (1985). Higher survival of an aposematic than of a cryptic form of a distasteful bug. *Oecologia, 67*, 411–415.

Sillén-Tullberg, B. (1988). Evolution of gregariousness in aposematic butterfly larvae: a phylogenetic analysis. *Evolution, 42*, 293–305.

Sillén-Tullberg, B., & Bryant, E. H. (1983). The evolution of aposematic coloration in distasteful prey: an individual selection model. *Evolution*, *37*, 993–1000.

Speed, M. P. (1993a). Muellerian mimicry and the psychology of predation. *Animal Behaviour, 45*, 571–580.

Speed, M. P. (1993b). When is mimicry good for predators? *Animal Behaviour, 46*, 1246–1248.

Speed, M. P. (1999). Robot predators in virtual ecologies: the importance of memory in mimicry studies. *Animal Behaviour, 57*, 203–213.

Speed, M. P. (2000). Warning signals, receiver psychology and predator memory. *Animal Behaviour, 60*, 269–278.

Speed, M. P. (2001a). Batesian, quasi-Batesian or Muellerian mimicry? Theory and data in mimicry research. *Evolutionary Ecology, 13*, 755–776.

Speed, M. P. (2001b). Can receiver psychology explain the evolution of aposematism? *Animal Behaviour*, *61*, 205–216.

Speed, M. P., & Ruxton, G. D. (2002). Animal behaviour: Evolution of suicidal signals. *Nature*, *416*, 375–377.

Speed, M. P., & Turner, J. R. G. (1999). Virtual predators, receiver psychology and doubts about Müllerian mimicry: comments on MacDougall & Dawkins. *Animal Behavior, 58*(4), F10–F13.

Srygley, R. B. (2004). The aerodynamic costs of warning signals in palatable mimetic butterflies and their distasteful models. *Proceedings of the Royal Society of London: Biological Sciences, 271*, 589–594.

Staddon, J. E. R., & Gendron, R. P. (1983). Optimal detection of cryptic prey may lead to predator switching. *American Naturalist, 122*(6), 843–848.

Struhsaker, T. T. (1967). Auditory communication among vervet monkeys. In Altmann, S. (Ed.), *Social Communication among Primates*. University of Chicago Press, Chicago.

Summers, K., & Clough, M. E. (2001). The evolution of coloration and toxicity in the poison frog family (Dendrobatidae). *Proceedings of the National Academy of Sciences of the United States of America, 98*, 6227–6232.

Swynnerton, C. F. M. (1915). A brief preliminary statement of a few of the results of five years' special testing of the theories of mimicry. *Proceedings of the Entomological Society of London, I*, xxxii–xliv.

Symula, R., Schulte, R., & Summers, K. (2001). Molecular phylogenetic evidence for a mimetic radiation in Peruvian poison frogs supports a Müllerian mimicry hypothesis. *Proceedings of the Royal Society of London*, *268*, 2415–2421.

Taylor, C., & Jefferson, D. (1994). Artificial life as a tool for biological inquiry. *Artificial Life*, *1(1/2)*, 1–13.

Thompson, J. N. (1994). *The Coevolutionary Process*. University of Chicago Press, Chicago IL.

Tinbergen, L. (1960). The natural control of insects in pine woods. I. Factors influencing the intensity of predation by songbirds. *Archives Neelandaises de Zoologie*, *12*, 265–343.

Turner, J. R. G. (1975). A tale of two butterflies. *Natural History*, *84*(2), 28–37.

Turner, J. R. G. (1977). Butterfly mimicry: The genetical evolution of an adaptation. *Evolutionary Biology*, *10*, 163–206.

Turner, J. R. G. (1981). Adaptation and evolution in *Heliconius*—a defense of neodarwinism. *Annual Review of Ecology and Systematics*, *12*, 99–121.

Turner, J. R. G. (1983). Mimetic butterflies and punctuated equilibria: some old light on a new paradigm. *The Biological Journal of the Linnean Society*, *20*, 277–300.

Turner, J. R. G. (1984). Mimicry: the palatability spectrum and its consequences. In Vane-Wright, R. I., & Ackery, P. R. (Eds.), *The Biology of Butterflies (Royal Entomological Society of London Symposium. No. 11)*, pp. 141–161. Academic Press, New York.

Turner, J. R. G. (1987). The evolutionary dynamics of Batesian and Mullerian mimicry: Similarities and differences. *Ecological Entomology*, *12*, 81–95.

Turner, J. R. G. (1988). The evolution of mimicry: A solution to the problem of punctuated equilibrium. *American Naturalist*, *131*, S42–S66.

Turner, J. R. G. (1995). Mimicry as a model for coevolution. In Arai, R., Kato, M., & Doi, Y. (Eds.), *Biodiversity and Evolution*, pp. 131–150. National Science Museum Foundation, Tokyo.

Turner, J. R. G., Kearney, E. P., & Exton, L. S. (1984). Mimicry and the Monte Carlo predator: The palatability spectrum and the origins of mimicry. *Biological Journal of the Linnean Society*, *23*, 247–268.

Turner, J. R. G., & Speed, M. P. (1996). Learning and memory in mimicry. I. Simulations of laboratory experiments. *Philosophical Transactions of the Royal Society of London: Biological Sciences*, *351*, 1157–1170.

Turner, J. R. G., & Speed, M. P. (1999a). How weird can mimicry get? *Evolutionary Ecology*, *13*, 807–827.

Turner, J. R. G., & Speed, M. P. (1999b). Learning and memory in mimicry. II. Do we understand the mimicry spectrum? *Biological Journal of the Linnean Society*, *67*, 281–312.

Van Valen, L. (1973). A new evolutionary law. *Evolutionary Theory*, *1*, 1–30.

Wallace, A. R. (1858). On the tendency of varieties to depart indefinitely from the original type. *Proceedings of the Linnean Society (Zoology)*, *3*, 53–62.

Wallace, A. R. (1866). On the phenomena of variation and geographical distribution, as illustrated by the papilionidae of the malayan region. *Transactions of the Linnean Society of London*, *35*, 1–72.

Wallace, A. R. (1867). S129 discussion..

Wallace, A. R. (1871). *Contributions to the theory of natural selection* (2nd edition). MacMillan and Co., London.

Wallace, A. R. (1879a). The protective colours of animals. *Science for All*, 138–157.

Wallace, A. R. (1879b). Protective mimicry in animals. *Science for All*, 284–296.

Wallace, A. R. (1889). *Darwinism—an exposition of the theory of natural selection with some of its applications*. MacMillan and Co, London.

Walther, F. R. (1969). Flight behaviour and avoidance of predators in the Thomson's gazelle (*Gazella thomsoni* Guenther 1884). *Behaviour*, *34*, 435–442.

Wasserthal, L. T. (1975). The role of butterfly wings in regulation of body temperature. *Journal of Insect Physiology*, *21*, 1921–1930.

Watson, R. A., & Pollack, J. B. (2001). Coevolutionary dynamics in a minimal substrate. In L. Spencer, e. a. (Ed.), *Proceedings of the 2001 Genetic and Evolutionary Computation Conference*. Morgan Kauffman.

Wickler, W. (1968). *Mimicry in Plants and Animals*. Wiedenfeld and Nicholson, London.

Wüster, W. (2004). Do aposematism and Batesian mimicry require bright colours? A test, using European viper markings. *Proceedings of the Royal Society of London: Biological Sciences*, *271*, 2495–2499.

Yachi, S., & Higashi, M. (1998). The evolution of warning signals. *Nature*, *394*, 882–884.

Yamauchi, A. (1993). A population-dynamic model of Batesian mimicry. *Researches on Population Ecology*, *35*, 295–315.

Zahavi, A. (1975). Mate selection—a selection for a handicap. *Journal of Theoretical Biology*, *53*, 205–214.

Zahavi, A. (1977). The cost of honesty (further remarks on the handicap principle). *Journal of Theoretical Biology*, *67*, 603–605.

Zahavi, A., & Zahavi, A. (1997). *The Handicap Principle: A Missing Piece of Darwin's Puzzle*. Oxford University Press, Oxford.