

University of Alberta

Exploring the dynamics of biological game theory models with genetic algorithms.

by

Steven Hamblin



**A thesis submitted to the Faculty of Graduate Studies and Research
in partial fulfillment of the requirements for the degree of**

Master of Science

Department of Psychology

Edmonton, Alberta

Fall, 2007



Library and
Archives Canada

Bibliothèque et
Archives Canada

Published Heritage
Branch

Direction du
Patrimoine de l'édition

395 Wellington Street
Ottawa ON K1A 0N4
Canada

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file *Votre référence*
ISBN: 978-0-494-33251-1
Our file *Notre référence*
ISBN: 978-0-494-33251-1

NOTICE:

The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

AVIS:

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protègent cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.


Canada

*Dedicated to Brandy, without whom
this thesis wouldn't have been worth writing.*

Abstract

Game theory models are widely used in biology to model behaviour, but the dynamics of models which exceed the simplest levels of complexity have not been widely explored. This is due, in part, to the difficulty inherent in providing analytical solutions when model complexity is increased to allow for more realistic descriptions of behaviour. I explored two game theoretical models used in the study of animal communication using genetic algorithms - a powerful heuristic search technique - and found alternative, unknown solutions to both models. As well, the use of a genetic algorithm on these models highlighted that the dynamics of a model may not lead a population to adopt the mathematically optimal solution. I suggest that authors of game theory models should take these considerations into account when designing models to demonstrate a theoretical point, and advocate genetic algorithms as a sensible method for solving complex game theory models.

Acknowledgements

Research funding by NSERC (Canada) Discovery Grant to P. Hurd. I would like to thank the members of the Hurd lab for help at various stages of preparing this thesis: Alisha Brown, Walter Espinoza, Reginia Yan, Maria Modanu, Adam Reddon, Amy Chee. Thanks also to Geoff Hollis for discussions on genetic algorithm structure. For Chapter 2, thanks go to ES Adams for mentioning, at the 2003 ABS conference, that his genetic algorithms did not converge on the ESS for the E85 game, and the following discussion. Reginia Yan and Alisha Brown contributed helpful comments on a final draft of this chapter.

Contents

1	Introduction	1
1.1	Models of communication	2
1.2	Population dynamics	3
1.3	Game complexity	5
1.4	Chapter Overview	6
2	The E85 game	10
2.1	Introduction	12
2.2	Methods	18
2.2.1	The Conventional Signalling Game	18
2.2.2	The Genetic Algorithm	19
2.2.3	Characterizing outcomes	22
2.3	Results	24
2.3.1	Experiment 1 — ESS fixation as a function of mutation rate and seed level	24
2.3.2	Experiment 2 — Characterization of non-ESS outcomes.	25
2.4	Discussion	28

Appendices	43
2.A Stability concepts	43
2.B Performance of the GA on other tasks	45
2.C Sensitivity of the results to methods and parameters	46
2.D Proof that **AAAA is an ES Set	49
3 The Sir Philip Sidney Game	53
3.1 Introduction	55
3.2 Methods	58
3.2.1 Johnstone & Grafen’s Sir Philip Sidney game	58
3.2.2 The Genetic Algorithm	61
3.2.3 Characterizing Outcomes	63
3.2.4 Parameter randomization	63
3.3 Results	64
3.3.1 Genetic algorithm simulations	64
3.3.2 Payoff analysis	65
3.3.3 Re-parameterization of the model	67
3.4 Discussion	70
4 Conclusion	74

List of Figures

2.1	A non-ESS extensive form game solution.	15
2.2	Representation of strategy change over time.	23
2.3	Possibility of ESS fixation.	26
2.4	Breakdown of outcomes in the TCNP model.	29
2.5	Breakdown of outcomes in the VCDF model.	30
3.1	Signalling and donation strategies.	56
3.2	SPS game chromosomes.	62
3.3	Simulation outcomes.	66

List of Tables

2.1	Encoding strategies onto chromosomes.	21
2.2	Colours used to map strategy evolution.	41
2.3	Breakdown of unclassified outcomes in the simulation results. . .	42
C.4	Effect of varying genetic algorithm parameters on the simulation outcomes.	48
D.5	Stability table for ABAAAA.	50
D.6	Stability table for ABAFAA.	51
D.7	Stability table for AAAAAA.	51
3.1	Algebraic payoff matrix for the SPS game.	60
3.2	Numerical payoff matrix for the SPS game.	61
3.3	Numerical payoffs to the semi-separating ESS and the ES Set at equilibrium.	67
3.4	Regions in the parameter space.	69

Chapter 1

Introduction

This thesis explores the population dynamics of game theory models of biological communication using genetic algorithms. Game theoretical models attempt to capture the essence of a social interaction by stripping them down to their barest minimum. These caricatures may not accurately depict biologically realistic situations, but they are required to be kept very simple in order to be formally solvable. I advocate genetic algorithms as a heuristic search technique for identifying potential solutions to game theory models which are too complex to completely solve analytically. I also attempt to demonstrate that the assumptions about evolutionary dynamics which have underlied the use of Evolutionarily Stable Strategy (ESS) analyses in biological game theory models are flawed, and suggest that they should be re-examined in the light of the dynamics which emerge from these genetic algorithm simulations.

1.1 Models of communication

The study of animal communication has been an important part of the study of animal behaviour (Tinbergen 1951; Dawkins and Krebs 1978; Dawkins 1982; Krebs and Dawkins 1984). Specifically, why would communication between animals ever evolve in the absence of strictly common interests? If animals have overlapping interests there is obviously no paradox, as the evolution of signalling then leads to increased benefits for both parties. Yet if the relationship between two animals includes conflicting interests, then communication suddenly becomes much more difficult to explain. Classic examples of such communication can be found in the study of animal aggression, where signalling seems entirely paradoxical: if an animal can gain an advantage during a conflict with another animal (conspecific or not) by producing a signal, such as "I am strong" (an example of a signal indicating Resource Holding Potential, or RHP), then it would seem to be to the benefit of all animals in this situation to produce the same signal, whether they are strong or not (Dawkins and Krebs 1979). But if all animals produce the same signal, then the signal loses its usefulness - i.e. it no longer says anything interesting to the other parties in the conflict, since even weak animals signal that they have high RHP - and it should quickly lose the power to arouse a response.

Many ideas have been advanced to deal with this question. Early efforts centred around cooperative or group selectionist arguments (reviewed in Dawkins and Krebs 1978; Krebs and Dawkins 1984), but these were challenged by the introduction of game theoretical modelling. Building on the work of W. D. Hamilton (Hamilton 1967) and George Price, John Maynard Smith (Maynard Smith and Price 1973; Maynard Smith 1974; 1982) pioneered the use of game theory in modelling biological situations by advancing the notion of the Evolutionarily Stable Strategy (ESS), and the technique was quickly applied to questions of animal communication (e.g.

Caryl 1979; Hinde 1981; Enquist 1985; Grafen 1990; Grafen and Johnstone 1993), and continues to be widely used to this day (Johnstone 1998; Bradbury and Vehrencamp 1998; Maynard Smith and Harper 2003; Hurd and Enquist 2005).

ESS theory revolves around questions of *evolutionary stability*. In intuitive terms, an ESS is a strategy which is “uninvadable” : if a population is composed entirely of individuals playing that strategy, there is no mutant which can invade in small proportions. ESS theory has led to a great shift in the way that many biological problems are thought about and modelled, due largely to its plausible and easily-understood nature. It was with game theory that models of handicapped (Zahavi 1975; 1977; Grafen 1990) and conventional (Enquist 1985; Hurd and Enquist 1998) signalling were first created, and it was with ESS theory that they were shown to be evolutionarily stable.

1.2 Population dynamics

It is here that I turn my focus from animal communication to the population dynamics of these game theory models.

Most analyses of game theory models in biology has been restricted to the hunt for the ESS (or ESSs) of the game. In fact, many of these sort of models are implicitly constructed with a simple process: after identifying the theoretical issue of interest, the authors of the models come up with a plausible game theory model that will contain the desired strategy and then conclude by demonstrating that this strategy is, in fact, evolutionarily stable. In itself, the procedure is reasonable and works well to highlight theoretical points.

However, in the process of generating these models, important questions regarding population dynamics are often left by the wayside. To begin, an ESS is guaranteed - by definition - to be stable against small invasions (by a single strategy), but one might ask how a population comes to be at an ESS. If it is not composed entirely of ESS players to begin with, is it reasonable to assume that the evolutionary process will drive the population to the ESS? In mathematical terms, this is the same as asking if an ESS is an attracting state, whether globally or locally.

The answer to that question is a resounding “no”, at least not always. Nowak (1990) was among the first to identify this problem, when he wrote:

If all members of a population use an evolutionarily stable strategy (Maynard Smith, 1972), then no mutant strategies can invade. However, it is a common but misleading conclusion that evolution will tend towards such a strategy (p. 237).

These words have not received their deserved attention since their publication (exceptions include Takada and Kigami 1991; Fogel et al. 1997;

Orzack and Hines 2005), and further problems abound. Models with complexity exceeding the most basic level of 2x2 matrix games are difficult to completely analyze by hand, and even the relatively implausible models such as those analyzed in this thesis push the limits of analytical approaches. Adding in more complexity by modelling biologically realistic states such as multiple signals, RHP states, temporal sequences, etc. can create models with millions of strategies, which are analytically intractable.

To help solve this problem, I have applied the powerful heuristic (trial-and-error) search tool known as a genetic algorithm to game theory models. Genetic algorithms simulate the processes of evolution, such as mutation, recombination, and selection, to evolve potential candidate solutions until an optimum is reached (Mitchell 1998). Genetic algorithms are widely used in many areas (Mitchell 1998), and their use has been suggested for game theory models before (Sumida et al. 1990), though the call has gone relatively unheeded (examples include Barta et al. 1997; Huse et al. 1999). To demonstrate the potential of genetic algorithms to solve game theory models which are too complex for analytical approaches, I applied a genetic algorithm to a well-known model of conventional signalling (Enquist 1985, hereafter referred to as the E85 model). The E85 game models an aggressive interaction between two players, where each player has a strength state (RHP) of “Strong” or “Weak” and can send one of two signals to indicate their strength state before choosing one of three actions. Because there is no cost to producing the signal, signal reliability is maintained by receiver response. For instance, a player signalling “Strong” may end up in an escalated fight that will end badly if they have chosen to signal unreliably. (Strictly speaking, because the signals are conventional, the form of the signal could be reversed without penalty. As long as strong signallers always use the “Weak” signal and vice versa, the model’s logic and the formal solutions are simply inverted.) The E85 model is among the most complex of models which currently exist in the animal communication literature and in the biological game theory literature in general. The results of applying a genetic algorithm to the E85 model, which are presented in Chapter 2, provide evidence that genetic algorithms are a sensible choice for analyzing hard game theory models.

1.3 Game complexity

In Chapter 3 I turned to a simpler game, intermediate in complexity between E85 and the overly simplistic realm of 2x2 matrix games. Sir Philip Sydney games (Maynard Smith 1991) are among the simplest published models of signalling, in which two players, the donor and beneficiary are each in one of two states (thirsty or not thirsty) and the beneficiary can choose one of two signals (also thirsty or not thirsty) to entice the donor to give a water resource to the beneficiary. Johnstone and Grafen (1993) used the Sir Philip

Sidney paradigm to create a model of handicapped signalling in which signalling only had to be “honest” on average - that is, free-riding by a distinct class of signallers could be tolerated at low levels if reliable signalling was to the benefit of the majority class. Johnstone and Grafen analyzed the model to demonstrate that this handicapped signalling strategy was an ESS, and so I used a genetic algorithm to explore the dynamics of this model. We found that the game’s simple form concealed a surprising fragility, since changing parameters to the model - while maintaining the conditions which guaranteed the original ESS - easily generated multiple solutions beyond those that had already been discovered.

1.4 Chapter Overview

Chapter 2 presents my work on the E85 model, in which I advocate the use of genetic algorithms and challenge the naive, ESS-focused analysis of game theory models. Chapter 3 is the report of our work on the simpler Sir Philip Sydney game of Johnstone and Grafen (1993), in which I continue to show that genetic algorithms are a good choice for solving game theory models and demonstrate that the problems I discussed in Chapter 2 apply to games of a surprisingly low level of complexity.

Bibliography

- Barta, Z., Flynn, R., and Giraldeau, L.-A. (1997). Geometry for a selfish foraging group: a genetic algorithm approach. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 264:1233–1238.
- Bradbury, J. W. and Vehrencamp, S. L. (1998). *Principles of animal communication*. Sinauer Associates, Inc., Massachusetts.
- Caryl, P. G. (1979). Communication by agonistic displays: what can games theory contribute to ethology? *Behaviour*, 68:136–169.
- Dawkins, R. (1982). *The extended phenotype: The long reach of the gene*. Oxford University Press, Oxford.
- Dawkins, R. and Krebs, J. R. (1978). Animal signals: Information or manipulation? In Krebs, J. R. and Davies, N. B., editors, *Behavioural ecology: An evolutionary approach*. Oxford, Blackwell, 1st edition.
- Dawkins, R. and Krebs, J. R. (1979). Arms races between and within species. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 205(1161):489–511.
- Enquist, M. (1985). Communication during aggressive interactions with particular reference to variation in choice of behaviour. *Animal Behaviour*, 33:1152–1161.
- Fogel, D. B., Fogel, G. B., and Andrews, P. C. (1997). On the instability of evolutionary stable strategies. *BioSystems*, 44:135–152.
- Grafen, A. (1990). Biological signals as handicaps. *Journal of Theoretical Biology*, 144:516–546.
- Grafen, A. and Johnstone, R. A. (1993). Why we need signalling theory. *Philosophical Transactions of the Royal Society, Biological Sciences*, 340:245–250.
- Hamilton, W. D. (1967). Extraordinary sex ratios. *Science*, 156(3774):477–488.
- Hinde, R. A. (1981). Animal signals: ethological and games-theory approaches are not incompatible. *Animal Behaviour*, 29:535–542.
- Hurd, P. L. and Enquist, M. (1998). Conventional signalling in aggressive interactions: the importance of temporal structure. *Journal of Theoretical Biology*, 192:197–211.
- Hurd, P. L. and Enquist, M. (2005). A strategic taxonomy of biological communication. *Animal Behaviour*, 70:1155–1170.

- Huse, G., Strand, E., and Giske, J. (1999). Implementing behaviour in individual-based models using neural networks and genetic algorithms. *Evolutionary Ecology*, 13:469–483.
- Johnstone, R. A. (1998). Game theory and communication. In Dugatkin, L. A. and Reeve, H. K., editors, *Game theory and animal behaviour.*, pages 94–117. Oxford University Press, New York.
- Johnstone, R. A. and Grafen, A. (1993). Dishonesty and the handicap principle. *Animal Behaviour*, 46:759–764.
- Krebs, J. R. and Dawkins, R. (1984). Animal signals: Mind-reading and manipulation. In Krebs, J. R. and Davies, N. B., editors, *Behavioural Ecology: An Evolutionary Approach*, chapter 15, pages 380–402. Blackwell Scientific, Oxford, 2nd edition.
- Maynard Smith, J. (1974). The theory of games and the evolution of animal conflicts. *Journal of Theoretical Biology*, 47(1):209–221.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, New York, NY.
- Maynard Smith, J. (1991). Honest signalling: the Sir Philip Sydney game. *Animal Behaviour*, 42:1034–1035.
- Maynard Smith, J. and Harper, D. (2003). *Animal Signals*. Oxford University Press, New York.
- Maynard Smith, J. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246(1):15–18.
- Mitchell, M. (1998). *An introduction to genetic algorithms*. MIT Press, Cambridge, Massachusetts, 1st paperback edition edition.
- Nowak, M. (1990). An evolutionarily stable strategy may be inaccessible. *Journal of Theoretical Biology*, 142:237–241.
- Orzack, S. H. and Hines, W. (2005). The evolution of strategy variation: Will an ess evolve? *Evolution*, 59(6):1183–1193.
- Sumida, B. H., Houston, A. I., McNamara, J. M., and Hamilton, W. D. (1990). Genetic algorithms and evolution. *Journal of Theoretical Biology*, 147:59–84.
- Takada, T. and Kigami, J. (1991). The dynamic attainability of ess in evolutionary games. *Journal of Mathematical Biology*, 29:513–529.
- Tinbergen, N. (1951). *The study of instinct*. Oxford University Press, London and New York.

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.

Chapter 2

The E85 game

Genetic algorithms and non-ESS solutions to game theory models.

Steven Hamblin

Department of Psychology
University of Alberta

2.1 Introduction

Game theory is one of the most well-established tools for modelling of social interactions, the application of which has led to notable advances in biological topics as diverse as sex ratio theory, cooperative behaviour, sexual selection, sperm competition, parent-offspring conflict, social foraging, and agonistic behaviour (Hamilton 1967; Maynard Smith 1982; Axelrod 1984; Andersson 1994; Giraldeau and Livoreil 1998; Reichert 1998). The topic of biological communication has been particularly influenced by game theoretical thinking (Johnstone 1997; 1998; Bradbury and Vehrencamp 1998; Maynard Smith and Harper 2003; Searcy and Nowicki 2005). The question of whether or not animals with conflicting interests should be expected to communicate reliably (or “honestly”) has been strongly shaped by both intuitive (e.g. Zahavi 1975; 1977; Dawkins and Krebs 1978; Caryl 1979; Hinde 1981; Krebs and Dawkins 1984) or formal (e.g. Enquist 1985; Enquist et al. 1985; Grafen 1990; Owens and Hartley 1991; Maynard Smith 1991; Johnstone and Grafen 1993; Johnstone and Norris 1993; Hurd 1995; Kim 1995; Hurd and Enquist 1998) game theoretical models.

In this paper we draw attention to properties of communication which make it a particularly difficult subject to model using formal game theory. In particular, we argue that the highly influential concept of an evolutionary stable strategy (ESS) is a poorly suited solution concept for models of communication, and discuss the need for the use of more powerful analytic solution concepts on such models. To aid in this task, we explore genetic algorithms as an alternative method of finding analytic outcomes to a relatively complicated model of communication. Communication requires a more complex model structure than many other biological game theoretical problems; while we will focus on models of biological communication to provide a concrete setting for our discussion, the general lessons and caveats will apply to any game theoretical model of comparable structure. Any game theoretical model in which players have hidden underlying states, or make more than one move per game, will have the level of complexity required for these ESS-presenting effects to present a problem (Selten 1975; 1983; 1988; Cressman 2003).

Most behavioural biologists are familiar with game theoretical models presented in their normal form, as payoff matrices. This form is appropriate in games where all players choose and play their strategies simultaneously (Gibbons 1992). Such interactions, with a single simultaneous behaviour, fail to capture the essence of communication, and most other social interactions. Communication requires individuals not only react to each other in a temporal sequence, but these individuals will most likely also vary in some underlying states such as differing RHP (Resource Holding Potential) states, varying resource valuations, or territorial ownership status which may be hidden from each other (Johnstone 1998; Maynard Smith and Harper 2003). Extensive form games are the appropriate game theoretical formulation of

such interactions as they allow strategies to be explicitly composed of several different possible moves made under different contingencies (Selten 1983; 1988; Cressman 2003; Hurd and Enquist 2005). Any game theory model which incorporates individual variation in ability or need, or variable moves made in response to the choices of other players, cannot properly be expressed in payoff matrix form alone.

The extra clarity inherent in the extensive form is especially obvious when the game is one of imperfect information, where one or more choices may be made by a player who is unaware of the outcome of decisions made by other players. Card games such as poker are the classic examples of games of this type. Each player knows the cards in their hand, but not those of other players (except in a probabilistic sense) while actions such as bets are common knowledge and provide information (even if ambiguous) to receivers about the signaller's hand. The parallels to biological communication are clear, but the discussion is not limited to communication games. For instance, a forager in a social foraging game who makes use of one conditional strategy when hungry and another when relatively satiated, or a sequence of life-history decisions in which the payoffs are not only frequency dependent but a function of the earlier and subsequent decisions, will have the same structural consequences as the communication game studied here.

Whether in their normal or extensive forms, game theoretical models in biology are typically solved by identifying evolutionary stable strategies (ESSs) (Parker 1984; Houston and McNamara 1999). ESSs are single strategies—or a single specific mixture of strategies (see Appendix 1)—which, if adopted by all members of a population, cannot be invaded by any other strategy (Maynard Smith 1982). The ESS concept has the virtue of intuitive clarity that allows for a greater understanding of social behaviour, but it is a mathematical formalism that is strongest when applied to simple models. Attempts to increase the biological realism of these models, by adding variables such as strength states or signalling stages, can easily create a model which has no ESS (Kim 1995). This problem is especially pressing in extensive form games (Selten 1975; 1980; 1983; 1988), due to two interrelated problems: non-pervasive strategies, and large strategy spaces.

The first, and most serious, problem in solving extensive form games stems from the formal conditions required to be an ESS (Appendix 1), which can only be met if the strategy is “pervasive” (Selten 1983). Pervasiveness is a property of strategies in which all possible decision points (information sets) in the game are reached with a non-zero probability. For example, consider the game presented in Figure 2.1. Player 2's node “b” is never reached when the players are playing optimally, and therefore neither are Player 1's nodes “c” and “d”. It does not matter which moves the players would make at these nodes, but the formal definition of an ESS requires not only that the moves to be made be specified, but that the strategy be stable against invasion by strategies that differ only in the moves at these nodes. This requirement that the “ESS” be stable against invasion by silent

mutations is obviously impossible to meet.

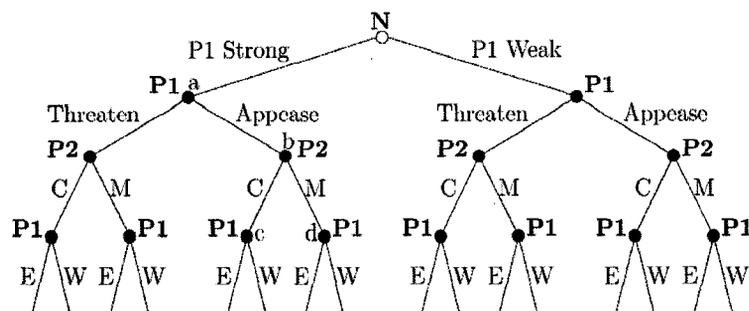


Figure 2.1: An example of a simple non-ESS extensive form game solution. The game has two players, Player 1 is either Strong or Weak, and chooses an initial move, either “Threaten” or “Appease”. Player 2 responds with either “Challenge” or “Mitigate”, then Player 1 moves again choosing one of “Engage” or “Withdraw”. Assume that evolution converges upon an optimal strategy in which Player 1 always chooses “Threaten” when Strong, so that the population evolves to a point at which, whenever the node marked “a” is reached, play never proceeds to the node marked “b”. Since “b” is never reached, Player 2’s behaviour at this node is never under any selective pressure, and neither is Player 1’s at the nodes “c” and “d”. Strategies which differ only in their choice of behaviour at these nodes are functionally equivalent, but no single strategy can possibly meet the definition of an ESS.

To cope with the problem of non-pervasive strategies, more powerful analytic solution concepts such as evolutionarily stable sets (ES Sets) must be used. In intuitive terms, an ES Set is a set of strategies which, individually, would be ESSs were it not for the fact that each member of the set scores equally well against all other members of the set, making the set’s members neutral to each other (Thomas 1985b;a; Cressman 1992; Balkenborg and Schlag 2001, see also Appendix 1). Thus, any strategy which would invade by a silent mutation will be part of the ES Set. Note that this is different from a mixed ESS in that a mixed ESS requires a precise mixture of pure strategies to be present, whereas for an ES Set, any of the members may be present in any ratio. If we view the players as being chosen from a polymorphic population where each member plays a pure strategy, then at

equilibrium a mixed ESS would be in the form of a specific ratio of each pure strategy in support of the ESS. An evolutionarily stable set would differ in that the equilibrium population could be equivalently composed of any conceivable mixture of the ES Set strategies, to the point where the entire population could play only one member of the set. ES Sets are not the only way to circumvent the pervasiveness issue (e.g. Limit ESSs: Selten 1983; 1988) but they are the most amenable to an intuitive definition and fit naturally with the results of our genetic algorithm explorations (see below).

The second problem to solving realistic communication games is that of large strategy spaces. Simple communication models are constrained to produce a single type of signal; more complicated models are required to produce phenomena such as conventional signals, which arguably have good empirical support (Hurd and Enquist 2001; Hurd 2004; Hurd and Enquist 2005). For example, in the simplest signalling game (Hurd 1995) each player has four pure strategies to choose between. The original Sir Philip Sidney game (Maynard Smith 1991) has four pure strategies per player, the version presented by Johnstone and Grafen (1993) has four for the donor and 16 for the beneficiary, while Maynard Smith's (1994) mutual signalling version has four beneficiary strategies and 64 donor strategies. In Kim's (1995) aggressiveness signalling game, each player has 32 pure strategies, while the conventional signalling game (Enquist 1985) (hereafter called the E85 game) each player has 324 pure strategies. Increasing the number of player states in the E85 game from just two possible values, to two different states with three levels each (while retaining the dichotomous signals and three end moves) results in over ten million pure strategies per player. Even if an ESS exists in such a large game, other forms of solutions - such as ES Sets or limit cycles in strategy space - will become more likely as the strategy space expands. The existence of any single form of solution does not preclude other, potentially more biologically relevant, forms. This suggests that analytic demonstrations of stable outcomes ought to be supplemented by searches for other regular patterns of strategy change.

A commonly used alternative for identifying solutions to games is an approach known as "evolutionary game theory" (Hammerstein 1998) - or "replicator dynamics" - in which a reasonable method of strategy replication and population dynamic is posited, and the space of all possible population compositions is mapped analytically (Taylor and Jonker 1978; Hofbauer and Sigmund 1988; Cressman 2003). Evolutionary game theory analysis may be tractable when there are few pure strategies. However, since each new strategy adds another dimension to the strategy space, analyzing dynamics in a dozen, or more, dimensions is not.

Further, in terms of population evolution, equilibria such as ESSs and ES Sets are strategies which are presupposed to have reached fixation in a population. There is nothing in the definition of an ESS which requires that selection take the strategy to fixation from even a small proportion of the phenotypic strategy space. In other words, while ESSs are stable points, they

are not necessarily attracting states (Nowak 1990). Games may also have more than one equilibrium solution, and the question of which solution has more attractive power in the strategy space may be analytically intractable. Games with large strategy spaces may even have non-equilibrium solutions, such as strategies with large attractive basins which are nonetheless beaten by an otherwise poor strategy, or quasi-stable outcomes such as limit cycles through strategy space. Although evolutionary game theory (replicator dynamics) can provide a picture of the population dynamics, we have already noted the problems of analyzing these dynamics in a large strategy space.

Analytic solutions are a way of describing what an evolutionarily stable population equilibrium will look like, in terms of what strategy or strategies will be used by the population when it is at that equilibrium. Evolutionarily stable sets and other solution concepts provide a remedy to the shortcomings of the ESS concept, but as noted in the introduction, a final problem remains: for all but the most simplistic models, finding these solutions can prove to be a vexing exercise. For games as large as E85 or the Kim game, considerable effort is required to find and demonstrate these solutions, and for games that are much larger than E85, the task may prove difficult or even impossible.

We suggest a possible alternative to deal with the computational challenge of finding solutions to more complex game theory models. Genetic algorithms (Holland 1975; Goldberg 1989; Sumida et al. 1990), computer simulations of strategy change by natural selection, hold great promise as a method for solving games. By tapping their heuristic optimization capabilities, a genetic algorithm can clarify attractive points in the strategy space which correspond to analytic solutions of game theory models, and by examining the evolution of the population over time and across runs, we can gain insight into the population dynamics. Despite the apparent promise, this approach remains relatively unexplored. To demonstrate the advantages of adopting this technique, we have applied a genetic algorithm to a well-studied game of communication, the conventional signalling game (Enquist 1985; Hurd 1997; Hurd and Enquist 1998; Enquist et al. 1998; Számadó 2000; 2003) in order to compare the simulation results to prior analytic work.

2.2 Methods

2.2.1 The Conventional Signalling Game

In the E85 game two players of varying discrete strength states compete over an indivisible resource, communicating with conventional signals before choosing behaviours with which to respond. The structure of the game is as follows:

1. Each player is randomly assigned a strength state, either “Strong” or “Weak” in a “move by nature”.

2. Knowing their own state, but not that of their opponent, each player then chooses a signal, (“A” or “B”) and sends it simultaneously.
3. Knowing their own state and the opponent’s signal, each player chooses and executes a behavioural response, one of “Full Attack”, “Pause-Attack”, or “Flee” (hereafter A, P or F) at the same time.
4. Payoffs are then calculated based on the players’ strength states and their behavioural responses.

Note that the choice of signal has no direct effect on payoff; any cost to making a particular signal will come from the response of the other player, which makes the signals “conventional” (Hurd and Enquist 2005). The E85 game, with its 324 pure strategies, is the simplest example of a conventional signalling game of which we are aware (Hurd and Enquist 2005).

The E85 model has more than one payoff representation. We used two different payoff quantifications for the genetic algorithms fitness function. The first, which we will refer to as TCNP, is the minimum variable version of the E85 game (Table A3.1 in Hurd and Enquist 1998), where $T = 1.0$, $C = 0.7$, $N = 0.4$, and $P = 0.1$. The second payoff scheme, referred to here as VCDF, is from the most biologically plausible version of the model (Table 1 in Hurd 1997); where $V = 100$, $C_{-1} = 15$, $C_0 = 15$, $C_1 = 70$, $F_a = 5$, $F_p = 5$.

Both the TCNP and VCDF payoffs used have only one ESS, the communicating strategy described by Enquist (1985) (actually, there are two ESSes, which are identical by symmetry, when the use of the costless signals is reversed; we can safely ignore this for the most part, but we will revisit the mirror ESS briefly in the results). The property of having a single ESS is useful in that it allows us to identify, a priori, a single strategy that should predominate. We can also use this strategy as a test of the genetic algorithm’s ability to find the “correct” solution.

2.2.2 The Genetic Algorithm

Simulations of the E85 model were done with a population size of 100 for a total of 500 generations. Each generation, we calculated the fitness of each population member as the mean from 5 plays of the game against other randomly chosen members of the population. Five percent of the population was killed each generation. The probability that a population member was killed was inversely proportional to their share of the total fitness of the population, and each killed member was replaced with a strategy chosen at random from the top 5% of the population, as ranked by fitness. After replacing the killed strategies, each locus of every member of the new population was mutated with a probability set by the mutation rate parameter. If a locus was selected to be mutated, the current value was

replaced equiprobably with any of the allowable values (including the original value) for that locus.

We tested the genetic algorithm on simpler problems with known solutions: the genetic algorithm performed well in finding the ESS to the Hawk-Dove game, and readily found the ESS to the E85 game when the opponent population was fixed to that strategy (see Appendix 2). Preliminary work with alternative genetic algorithm methods, such as other methods of selecting strategies to die or reproduce, or implementation of cross-over matings showed no discernible sensitivity to the details of the simulation methodology (see Appendix 3).

Representation of a strategy from the E85 game in the genetic algorithm was done using a six character chromosome in string form. Each locus specified the pure local strategy (*sensu* Selten 1983) for a different information set. All information sets in the game except for the move by nature which assigns strength states to each player corresponded to a chromosome locus, so that the entire chromosome codes for a pure behaviour strategy (*sensu* Selten 1983) with the noted exceptions. The first two loci represented the signalling portion of the chromosome, and the final four loci represented the end-game behaviours when: 1) ego state is strong and the opponent uses signal "A", 2) state is weak and opponent uses signal "A", 3) state is strong and opponent signals "B", and 4) when state is weak and opponent signals "B". The representation of the E85 ESS under this scheme is 'ABAFPA' - signal "A" when strong, signal "B" when weak, and end-game behaviours of attack, pause attack, flee, and pause attack under the combinations given above. Another example: a non-signalling strategy that displays "A" when strong and "A" when weak and then attacks at all end-move combinations would be represented as 'AAAAAA'. Further example strategies are presented in Table 2.1.

	Signal choice loci		End-game behaviour choice loci				Chromosome
	When Ego strong	When Ego weak	Opp. Signal is 'A'		Opp. Signal is 'B'		
			Ego strong	Ego weak	Ego strong	Ego weak	
locus number	1	2	3	4	5	6	
Example Strategies							
ESS	A	B	A	F	P	A	ABAFPA
Anti-ESS	B	A	P	A	A	F	BAPAAF
Non-signalling (A), all-attack	A	A	A	A	A	A	AAAAAA
Signalling, all- attack	B	A	A	A	A	A	BAAAAA
Non-Signalling (B), attack if strong, flee if weak	B	B	A	F	A	F	BBAFAF

Table 2.1: Encoding strategies onto chromosomes. Five example strategies are presented here as encoded into chromosome form for the genetic algorithm. The first two loci code for signal choice, either signal “A” or “B” when strong and weak respectively. If the signals used in these two strength states are the same, i.e. the alleles at loci number 1 and 2 are identical, then we say that the strategy is non-signalling. Loci 3 through 6 code for end-game moves as a function of the opponents signal choice and ego’s strength state.

We performed two sets of simulations. The first varied mutation rate from 0.001 to 0.005 (0.1 to 0.5%) at levels of initial “seed” (percentage of the population initialized to the ESS strategy) from 0 to 75%; ten runs were done at each combination of mutation rate and initial seed, for a total of 750 distinct simulations. This set was used to explore the probability of the ESS achieving fixation at different levels of the parameters (Experiment 1). The second set was used to investigate non-ESS endpoints found by the genetic algorithm (Experiment 2). These runs varied mutation rate from 0.001 to 0.005 at three low levels of initial seed (0, 5, and 10%) to explore what solutions the genetic algorithm would find when starting away from a population composed largely of ESS players. One hundred runs at each combination of mutation rate and seed were done for this set, giving a total of 1500 distinct simulations. This set was used to enumerate the various outcomes of the genetic algorithm and determine the formal solutions that arose from the runs.

2.2.3 Characterizing outcomes

For each simulation we plotted the change in strategy composition of the population over the course of the 500 generations. To plot the strategies we first divided each into two components, those loci used when weak, and those used when strong. Each of these two components had 18 possible permutations and can be represented as a point in a three dimensional array. One dimension indicated which signal the player uses (A or B), the other two dimensions indicated which behaviour to use (A, P or F) in response to opponent signal A, and which behaviour to use (again, A, P or F) in response to opponent signal B. Nine colours were chosen to represent the possible allelic combinations of the two behaviour loci (Table 2.2). We used the red-green colour dimension for the “Always Attack” to “Always Flee” dimension, and the yellow-blue colour dimension for the orthogonal dimension (running from Flee from A signallers and Attack B signallers —yellow— to Attack A signallers and Flee B signallers —blue). Dark or light shades of these colours indicated whether the player’s signal move was A or B respectively.

This colour coding was used to graph the evolution of strategies over the 500 generations (Figure 2.2). The strategy each run fixated on was recorded, or if no strategy predominated, this was also recorded. A strategy was considered to have gone to fixation if it had achieved 80% representation in the population for the last 50 generations (10% of the simulation). 80% dominance was chosen to serve as a criterion for those runs with higher mutation rates, where mutation pressure could create a large amount of noise variation from the clearly dominant strategy. Only a small number of cases required this liberal threshold. For example, in the simulations done for Experiment 2, only 87 (7%) of the 1263 runs that were labeled as having achieved fixation did so at less than 95% representation.

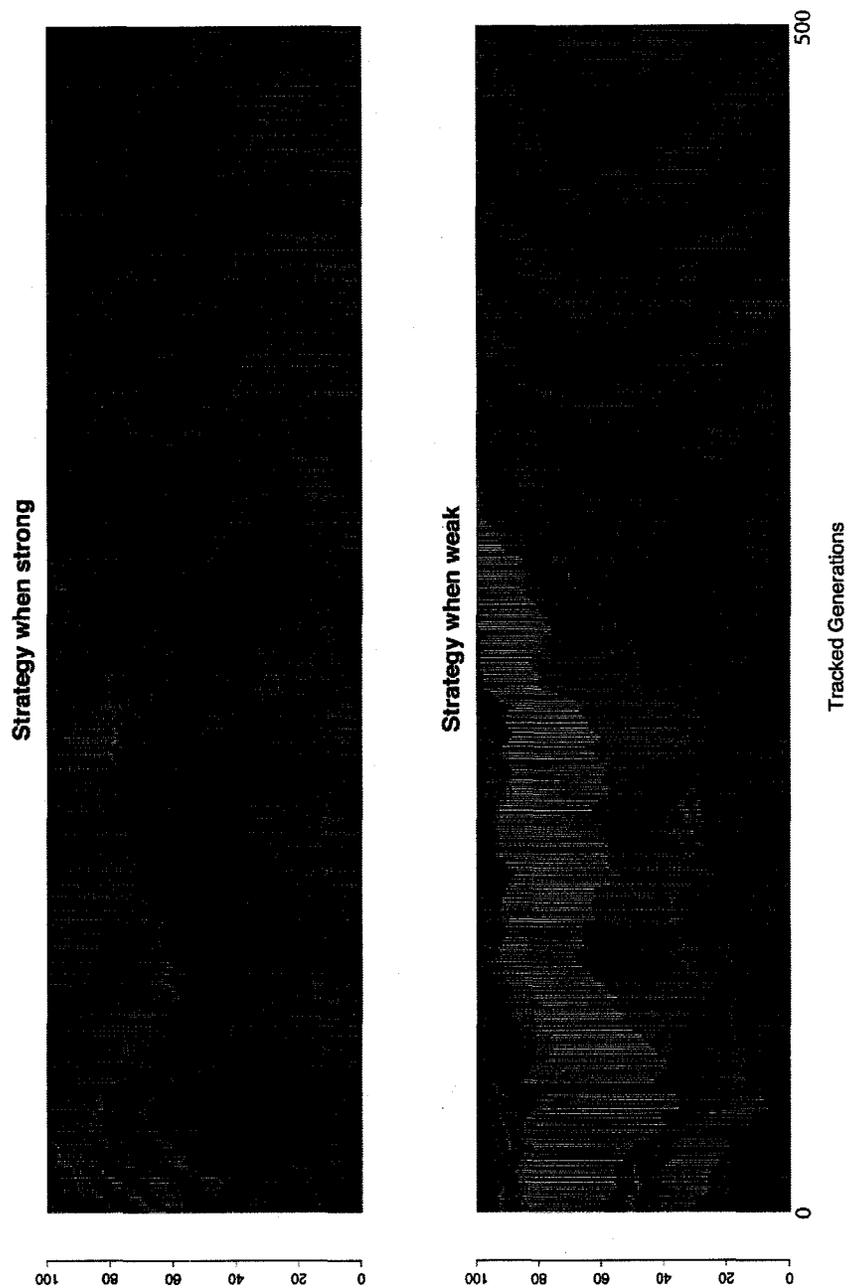


Figure 2.2: A representative graph of strategy change over 500 discrete generations of time. The run depicted was conducted at mutation rate 0.003 and initial ESS (light purple — dark yellow) seed of 10%. The ES Set (pink — red) reaches fixation by approximately the 300th generation.

2.3 Results

We tested the performance of the GA in a number of runs in which the opponent population was held fixed at the ESS. Instead of playing the game against other random members of the population, strategies played against only the ESS. In every single run, the GA evolved from a random starting population to fixation on the ESS in under 100 generations.

2.3.1 Experiment 1 — ESS fixation as a function of mutation rate and seed level

The genetic algorithm only converged on the ESS when the initial population was seeded with a high proportion of ESS players (Figure 3.3). Both payoff parameters and mutation rate influenced the probability that the ESS would go to fixation. Fixation on the ESS was more likely in the TCNP payoff version of the model. At least 10 to 15% initial ESS players was required to ensure a reliable fixation on the ESS in this case. Mutation rate had a larger effect on probability of ESS fixation in the VCDF payoff version.

Progressively higher rates made it more and more difficult for the genetic algorithm to fixate on the ESS. Even at low mutation rates, 20–30% of the initial population had to be playing the ESS before fixation was likely. In both cases, the ESS is clearly a very weak attractor, the probability that a randomly generated population will evolve to the ESS is virtually nil, particularly at high mutation rates.

2.3.2 Experiment 2 — Characterization of non-ESS outcomes.

While the ESS appears to have a very small basin of attraction (only 11% of the runs converged on the ESS; Figs. 2.4 & 2.5), there is another outcome which the genetic algorithm converges upon with high probability for most parameter combinations (73% of the runs). This outcome corresponds to the strategy set **AAAA, where * is any choice of allele at the loci responsible for choice of signal. None of these strategies is an ESS because they are invaded by their fellow set members, but the set as a whole meets the definition of an evolutionarily stable set (Thomas 1985b;a, Appendix 4). This solution has not been identified in previous analytical work on the model (e.g. Enquist 1985; Hurd 1997; Hurd and Enquist 1998; Számadó 2000; 2003), and will hereafter be referred to as the All-Attack ES Set.

We can easily verify that the ESS is a global optimum and the ES Set is a local optimum by examining the payoffs to ESS and ES Set players in populations of either ESS or ES Set players. From the stability table (see Appendix 2 in Hurd and Enquist 1998), we can show that the expected payoff in a population of ESS players (summed over all four possible ego and opponent states) is:

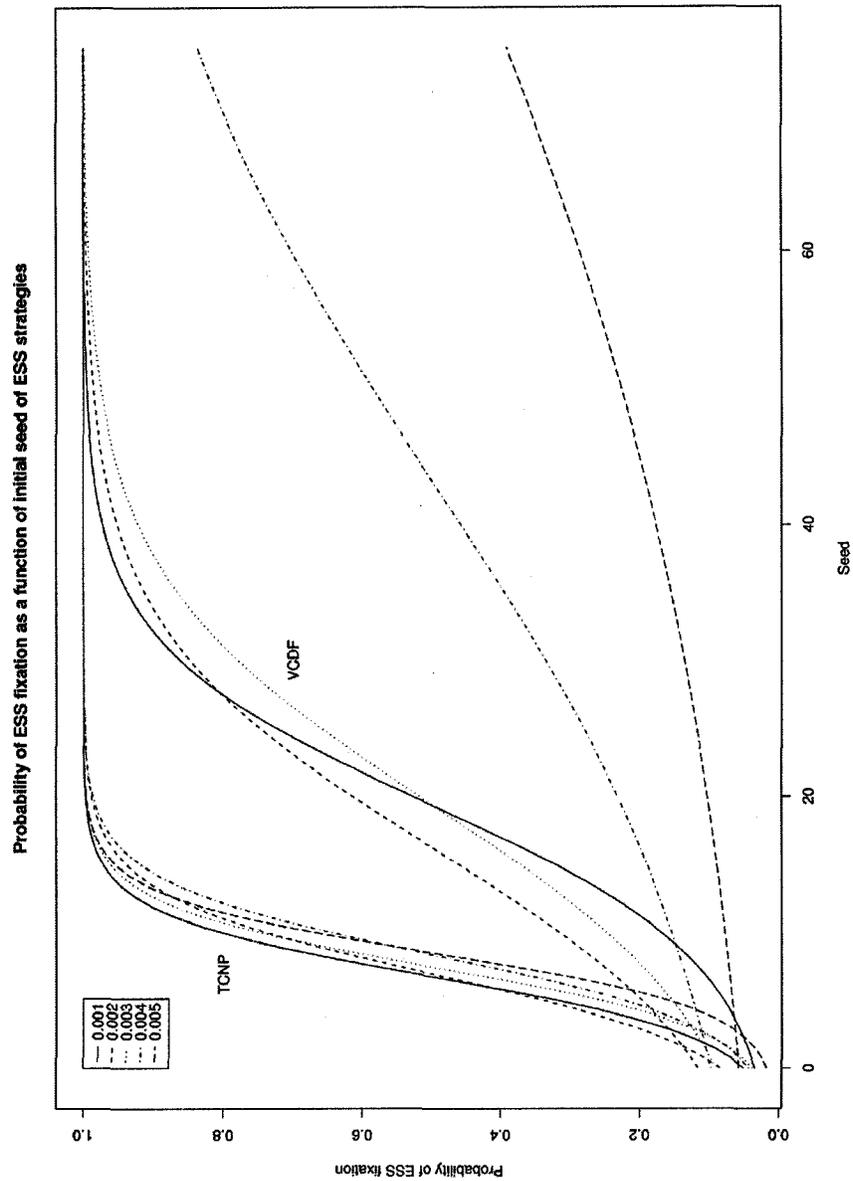


Figure 2.3: The probability that the ESS goes to fixation as a function of initial seed from 0 to 75%, increasing by 5% each time, and mutation rate from 0.001 to 0.005, increasing by 0.001 each time. Logistic curves were fit to data from 10 simulations per mutation rate and initial seed combination, for a total of 750 individual runs used.

$$= \left(\frac{1}{2}V - C_0\right) + V + 0 + \left(\frac{1}{2}V - C_0\right) \quad (2.1)$$

$$= 2V - 2C_0 \quad (2.2)$$

or, the value of winning twice minus the cost of fighting an opponent of equal strength twice.

The expected payoff to any member of the ES set playing in a population of ES Set players (from Table C.4 in Appendix 4) is:

$$= \left(\frac{1}{2}V - C_0\right) + (V - C_{-1}) + (-C_1) + \left(\frac{1}{2}V - C_0\right) \quad (2.3)$$

$$= 2V - 2C_0 - C_{-1} - C_1 \quad (2.4)$$

or, the value of winning twice minus the cost of fighting twice against an opponent of equal strength opponent and once each against a stronger and a weaker opponent.

The ESS equilibrium payoff is greater, on average, than the All-attack ES Set equilibrium payoff by $C_{-1} + C_1$, or the cost of two fights, one against a stronger opponent, the other against a weaker opponent, for every four encounters. From this, it is clear that a population composed of ESS players will be globally maximizing their fitness, making the ES set a local optimum.

The ESS and “Light red/Red” All-Attack evolutionarily stable set are not the only outcomes from the genetic algorithm simulations (Figs. 2.4 & 2.5). Two hundred and thirty six of the 1500 runs (16%) resulted in outcomes that were neither of the two stable outcomes identified above. These are grouped into several different classes summarized in Table D.7. In the first case, the genetic algorithm did not fixate on a single strategy, or set of strategies, and the population was still in flux at the end of the run. The second case is where the genetic algorithm was clearly trending toward either the ESS or ES Set but simply did not reach the fixation criteria in time. In some runs at higher mutation rates, this occurred when mutation pressure kept enough other players in the population to forestall any strategy from reaching the cutoff for fixation no matter how dominant it was. Between them, the strategies in the first two categories accounted for 12% of the 1500 runs of the VCDF model (see the first entry in Table D.6). The third case is comprised of outcomes where the genetic algorithm found the reversed signal convention version of the E85 ESS, the strategy BAPAAF. The small number of outcomes in which this occurred is unsurprising given that when a run was seeded with ESS players, it was done solely with the ABAFPA version of the ESS.

The fourth case includes all of those runs wherein the genetic algorithm fixated on a strategy or strategy set not yet mentioned. The single strategy outcomes (lines 2 through 8 in Table 2.3) are composed of non-signalling

strategies which appear to be indifferent at two loci and play Pause Attack at loci 3 (strategies which always signal A) or 4 (strategies which always signal B) as well as similarly playing Flee at loci 5 or 6. Analytical investigation has confirmed that these strategies are indifferent at some loci, similar to the All-Attack ES Set, but they do not meet the formal criteria of an evolutionarily stable set and should not be stable. Their ability to invade each other likely gives them a slightly greater attractive power in the strategy space, but they still amount to a relatively small basin of attraction that the genetic algorithm was stranded in by the end of the simulation. The sporadic appearance of these strategies and their lack of analytically demonstrated stability suggests that runs in which they appeared to be dominant were simply not continued long enough for them to be displaced. These cases do, however, highlight that the genetic algorithm will find outcomes that might be attracting, if only temporarily, even if they do not demonstrate analytical stability. They also highlight the continued need for analytical investigation to distinguish between and describe outcomes that are stable as opposed to those that merely appear stable.

In the fifth case are outcomes where the genetic algorithm appeared to fixate on a mixture of three or more strategies. These were not analysed extensively, but their low rate of occurrence (1.3%) suggests that these are unstable local attractors, though a more detailed analytical examination would be required to prove that conclusively.

2.4 Discussion

In our genetic algorithm investigation of the E85 conventional signalling game, we discovered a previously unknown ES Set solution to the game. The All-Attack ES Set is a non-communicating strategy—it does not respond to signals—and has a very large basin of attraction in the fitness landscape. This equilibrium is a local optima, it scores less at fixation than does the ESS.

What these results mean for the E85 game as a model of the evolution of conventional signalling is not clear. The simulations show the importance of assumptions about the evolutionary starting points of strategic signalling systems. The process of display ritualization (Tinbergen 1952) may mean that the strategic evolution of signal use starts so close to the eventual ESS that arguments based on the relative sizes of the attractive basins are moot. However, the difficulty in reaching the global optimum, the ESS, in this case is quite remarkable.

Since its formal introduction by Maynard Smith and Price (1973), the notion of an evolutionary stable strategy has had a profound impact not merely on formal models of social behaviour, but more informally in verbal models and discussions. Among non-theoreticians, the term “ESS” has become nearly synonymous with frequency dependence and game theoretical

Behaviour when:

Opp. Sig. is B	Opponent signal is A		
	A	P	F
A	Red	Pink	Yellow
P	Purple	Grey	Brown
F	Blue	Cyan	Green

Table 2.2: Colours used to map strategy evolution. Each of the nine possible behaviour strategies (for each ego strength state) was assigned a different colour. The Red-Green dimension was used to map the always Attack to always Flee dimension, and the Blue-Yellow dimension was used to map the other dimension. Colour saturation was used to map the choice of signal. A lighter shade denoted use of signal “A” and a darker shade the use of the signal “B”. This scheme allowed for 18 colours to represent behaviour in a given strength state, and therefore each of the 324 possible strategies could be denoted by a colour pair. The ESS is light purple – dark yellow.

Outcome	Mutation Rate					Total
	0.001	0.002	0.003	0.004	0.005	
Unstable mixture	21	25	27	40	36	149
Did not fixate in time	3	1	3	10	14	31
BAPAAF (alternative ESS)	8	8	6	0	0	22
BBAPAF	5	0	0	0	0	5
BBAPPF	2	0	0	0	0	2
AAPAAF	1	0	0	0	0	1
AAPPF	1	2	0	0	0	3
AAPAFF	0	1	0	0	0	1
BBAPPF	0	1	1	0	0	2
Stable mix of 3 or more strategies	11	3	6	0	0	20

Table 2.3: Breakdown of every unclassified outcome in 1500 runs of the VCDF model by mutation rate and outcome type. Note that stable non-ESS and non-ES set runs are increasingly uncommon as the mutation rate increases, with most non-optimal outcomes at mutation rates of 0.004 and 0.005 comprised of unstable mixtures or runs that did not reach the fixation criteria by the end of 500 generations. With the exception of the symmetric version of the ESS (BAPAAF), none of the alternative outcomes are analytically stable.

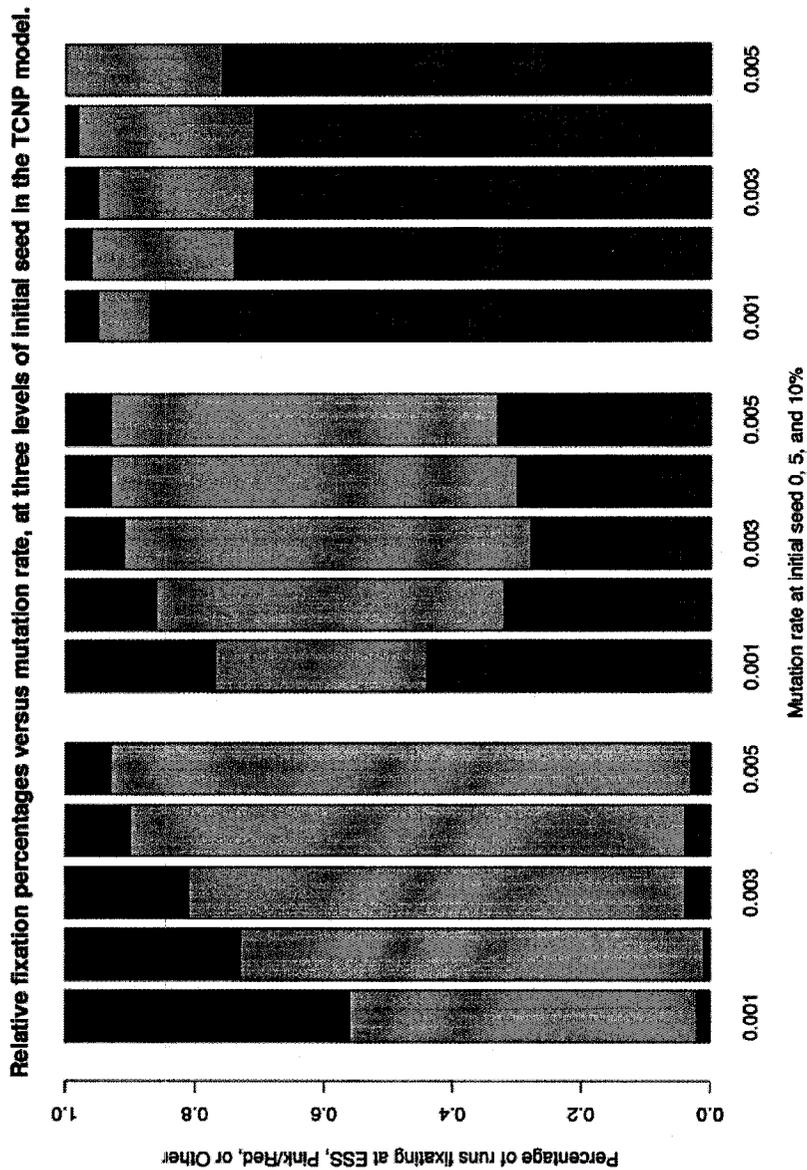


Figure 2.4: Stacked bar graph showing the percentage of runs for the TCNP model in which the ESS, ES Set, or Other strategies fixated at each level of mutation rate (0.001-0.005) and initial seed (0, 5, or 10). Each combination of mutation rate and initial seed was simulated 100 times, for a total of 1500 individual runs.

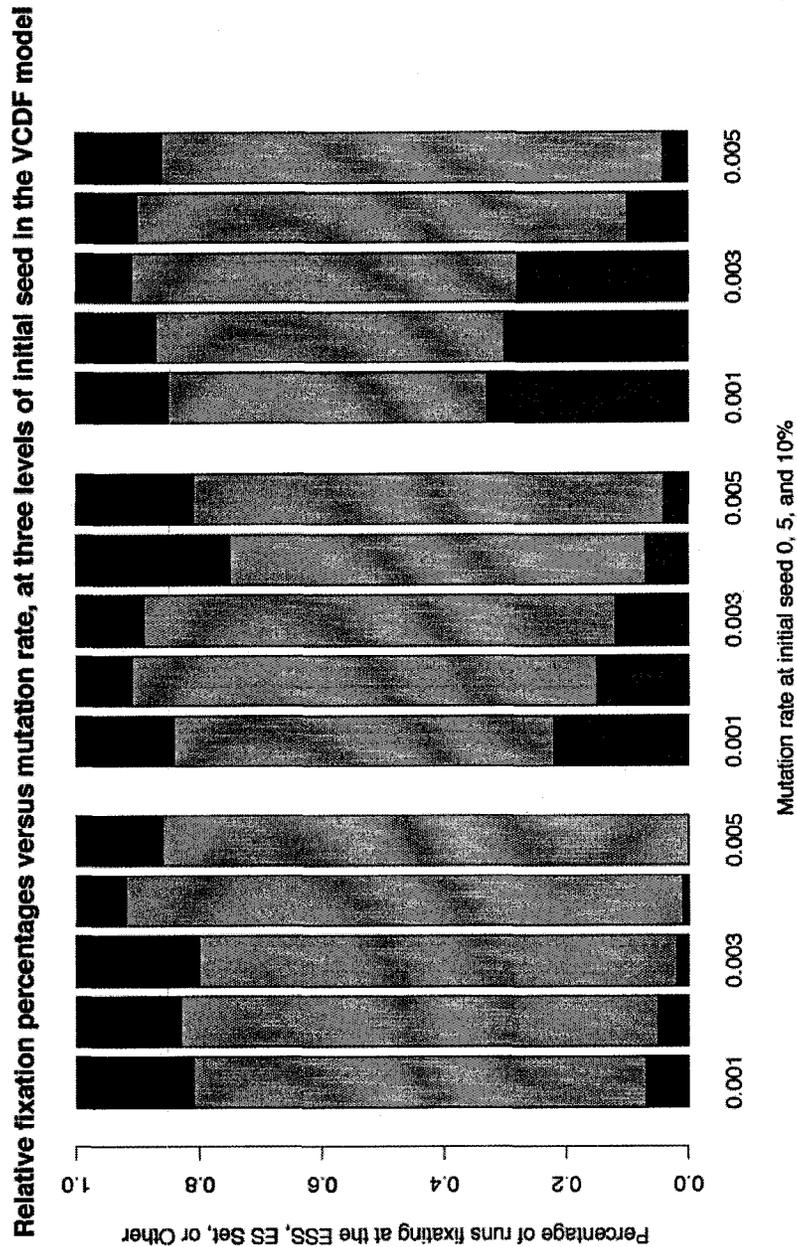


Figure 2.5: Stacked bar graph showing the percentage of runs for the VCDF model in which the ESS, ES Set, or Other strategies fixated at each level of mutation rate (0.001-0.005) and initial seed (0, 5, or 10). Each combination of mutation rate and initial seed was simulated 100 times, for a total of 1500 individual runs.

thinking in general. This term communicates an important and intuitive biological concept. While undoubtedly useful, this wider non-technical use risks theoretical work by focussing formal attention on this one single equilibrium definition. There are several reasons why it is necessary to widen the definition of evolutionary stability when solving communication models, and other games of similar complexity.

First, as in the E85 game studied here, ESSs may not be strong attractors. As Nowak (1990) noted of ESSs some time ago, "it is a common but misleading conclusion that evolution will tend towards such a strategy". Although other investigations have examined the dynamic stability of a population that is already at or in the neighbourhood of an evolutionarily stable solution (Nowak 1990; Takada and Kigami 1991; Thomas 1984; 1985b;a; Selten 1975; Leimar 1997; Taylor and Jonker 1978; Selten 1983; 1988; Hofbauer and Sigmund 1988), little attention has been paid to the process by which a population might come to approach an ESS in the first place. We have shown that a local optima in the form of an evolutionarily stable set with a large basin of attraction is the most likely evolutionary endpoint from most starting populations of this game.

Secondly, and possibly more importantly, is the likelihood that games will not have an ESS at all. Any game of realistic complexity - be they foraging games with hidden states influencing strategy choice, or games modelling any social interaction with more than one move per player - will likely have reasonable non-pervasive strategies. These games may have no ESSes at all, or may have ES Sets that are equally sound solutions as ESS. The number of strategies per player increases geometrically with increases in the complexity of extensive form games. While the condition of pervasiveness is met by the ESS solutions of communication games such as E85 and Kim's aggressiveness signalling game (Enquist 1985; Kim 1995), it becomes less and less likely as variables are added. Games any more complicated than E85, especially games with both variable RHP and subjective resource valuation states, are virtually guaranteed not to have ESSs for this reason. Methods using other solution definitions must be developed to allow for the complexity of biological models. Given that these models are exceptionally difficult to solve by hand, genetic algorithms provide one tractable alternative to analytical investigation. Genetic algorithms can never replace analytic work completely, even for complicated game theoretical problems. Ideally, once potential solutions are identified by the genetic algorithm, analytic methods may be brought to bear to investigate strategies identified by the genetic algorithm. This is the process that we followed in proving analytically that the solution the genetic algorithm converged on is an ES Set (see Appendix 4).

A potential caveat to the use of genetic algorithms on game theory models lies in the nature of the search space. A genetic algorithm is a powerful optimization tool with which we can quickly and reliably simulate the dynamics in games like the E85 game to determine what solution a random starting population will evolve to. However, despite the widespread use of

genetic algorithms to solve optimization problems in a variety of fields (Mitchell 1998), genetic algorithms work best on a solution space that is static and unchanging (e.g. a system of equations); less is known about the performance of genetic algorithms in a dynamic solution space, such as a constantly evolving population. The magnitude of this effect can be seen in the difference in fixation times when the target population is forced to remain at the ESS (< 100 generations) versus the time to fixation when the target population evolves (\approx 300 generations).

Previous investigations using genetic algorithms have usually either investigated the effect of stochastic events on very simple 2×2 matrix games (e.g. Maynard Smith 1988; Bergstrom and Godfrey-Smith 1998; Orzack and Hines 2005), or explored interactions far too complicated to be solved analytically (Barta et al. 1997; Hoffmeister and Roitberg 1998; Just and Morris 2003). More work needs to be done on games of intermediate complexity if genetic algorithms are to be applied to game theory with the intention of finding the solutions that would be found using analytic techniques. The results of the two techniques must be compared before their results are considered interchangeable. This caveat does not apply to the use of genetic algorithms in evolving neural networks to address the same issues (e.g. Enquist and Arak 1993; 1994; Johnstone 1994; Ezoe and Iwasa 1997; Blumstein et al. 2006; Huse et al. 1999). Neural network models seek to avoid assumptions made by game theoretical models about signal space divisions, and whose purpose is to produce results at odds with game theory. When stochastic simulations disagree with analytical results (e.g. McNamara et al. 2004), it may be argued that the simulations are more biologically important than the analytic results. In the present case, however the simulations identify an overlooked class of analytic solutions, ES Sets. These non-ESS evolutionarily stable outcomes which may, or may not, be more biologically plausible outcomes than the ESS depending on the assumptions made about the evolutionary starting point of the population.

Taken together, the questions raised by our application of evolutionary algorithms to the conventional signalling model present an interesting avenue for exploration in game theoretical approaches to evolutionary questions, as well as casting healthy skepticism on the biological relevance of ESSs. Not all evolutionarily stable outcomes are ESSs. ESSs are not the best, or only, solutions to biological game theory questions.

Bibliography

- Andersson, M. (1994). *Sexual selection*. Princeton University Press, Princeton, Massachusetts.
- Axelrod, R. (1984). *The evolution of cooperation*. Basic Books, New York, NY.
- Balkenborg, D. and Schlag, K. H. (2001). Evolutionarily stable sets. *International Journal of Game Theory*, 29:571–595.
- Barta, Z., Flynn, R., and Giraldeau, L.-A. (1997). Geometry for a selfish foraging group: a genetic algorithm approach. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 264:1233–1238.
- Bergstrom, C. T. and Godfrey-Smith, P. (1998). On the evolution of behavioral heterogeneity in individuals and populations. *Biology and Philosophy*, 13:205–231.
- Blumstein, D. T., Bitton, A., and DaVelga, J. (2006). How does the presence of predators influence the persistence of antipredator behavior? *Journal of Theoretical Biology*, 239:460–468.
- Bradbury, J. W. and Vehrencamp, S. L. (1998). *Principles of animal communication*. Sinauer, Sunderland, MA.
- Caryl, P. G. (1979). Communication by agonistic displays: what can games theory contribute to ethology? *Behaviour*, 68:136–169.
- Cressman, R. (1992). Evolutionary stable set in symmetric extensive two-person games. *Math. Biosci.*, 108:179–201.
- Cressman, R. (2003). *Evolutionary dynamics and extensive form games*. MIT Press.
- Dawkins, R. and Krebs, J. R. (1978). Animal signals: information or manipulation? In Krebs, J. R. and Davies, N. B., editors, *Behavioural Ecology: an evolutionary approach*, pages 282–309. Blackwell, Oxford, first edition.
- Enquist, M. (1985). Communication during aggressive interactions with particular reference to variation in choice of behaviour. *Anim. Behav.*, 33:1152–1161.
- Enquist, M. and Arak, A. (1993). Selection of exaggerated male traits by female aesthetic senses. *Nature*, 361:446–448.
- Enquist, M. and Arak, A. (1994). Symmetry, beauty and evolution. *Nature*, 372:172–175.

- Enquist, M., Ghirlanda, S., and Hurd, P. L. (1998). Discrete conventional signalling of continuously varying resource value. *Anim. Behav.*, 56:749–753.
- Enquist, M., Plane, E., and Roed, J. (1985). Aggressive communication in fulmars (*fulmarus glacialis*). *Anim. Behav.*, 33:1007–1020.
- Ezoe, H. and Iwasa, Y. (1997). Evolution of condition-dependent dispersal: A genetic-algorithm search for the ESS reaction norm. *Researches on Population Ecology*, 39:127–137.
- Gibbons, R. (1992). *A primer in game theory*. Harvester Wheatsheaf, London.
- Giraldeau, L.-A. and Livoreil, B. (1998). Game theory and social foraging. In Dugatkin, L. A. and Reeve, H. K., editors, *Game theory and animal behaviour.*, pages 16–37. Oxford University Press, New York.
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley, Reading, Massachusetts.
- Grafen, A. (1990). Biological signals as handicaps. *J. theor. Biol.*, 144:517–546.
- Hamilton, W. D. (1967). Extraordinary sex ratios. *Science*, 156(3774):477–488.
- Hammerstein, P. (1998). What is evolutionary game theory? In Dugatkin, L. A. and Reeve, H. K., editors, *Game theory and animal behaviour.*, pages 3–15. Oxford University Press, New York.
- Harsanyi, J. (1973). Oddness of the number of equilibrium points: a new proof. *Int. J. Game Theory*, 2:235–250.
- Hinde, R. A. (1981). Animal signals: ethological and games-theory approaches are not incompatible. *Anim. Behav.*, 29:535–542.
- Hofbauer, J. and Sigmund, K. (1988). *Evolutionary Games and Population Dynamics*. Cambridge U. Press, Cambridge.
- Hoffmeister, T. S. and Roitberg, B. D. (1998). Evolution of signal persistence under predator exploitation. *Ecoscience*, 5:312–320.
- Holland, J. H. (1975). *Adaptation in natural and artificial systems*. University of Michigan Press.
- Houston, A. I. and McNamara, J. M. (1999). *Models of adaptive behaviour: an approach based on state*. Cambridge University Press, Cambridge, Massachusetts.

- Hurd, P. L. (1995). Communication in discrete action–response games. *J. theor. Biol.*, 174:217–222.
- Hurd, P. L. (1997). Is signalling of fighting ability costlier for weaker individuals? *J. theor. Biol.*, 184:83–88.
- Hurd, P. L. (2004). Conventional displays: evidence for socially mediated costs of threat displays in a lizard. *Aggressive Behavior*, 30:326–431.
- Hurd, P. L. (2006). Resource holding potential, subjective resource value, and game theoretical models of aggressiveness signalling. *Journal of Theoretical Biology*, 241:639–648.
- Hurd, P. L. and Enquist, M. (1998). Conventional signalling in aggressive interactions: the importance of temporal structure. *J. theor. Biol.*, 192:197–211.
- Hurd, P. L. and Enquist, M. (2001). Threat display in birds. *Can. J. Zool.*, 79:931–942.
- Hurd, P. L. and Enquist, M. (2005). A strategic taxonomy of biological communication. *Anim. Behav.*, 70:1155–1170.
- Huse, G., Strand, E., and Giske, J. (1999). Implementing behaviour in individual-based models using neural networks and genetic algorithms. *Evolutionary Ecology*, 13:469–483.
- Johnstone, R. A. (1994). Female preference for symmetrical males as a by-product of selection for mate recognition. *Nature*, 372:172–175.
- Johnstone, R. A. (1997). The evolution of animal signals. In *Behavioural Ecology: an evolutionary approach*, pages 155–178. Blackwell, Oxford, fourth edition.
- Johnstone, R. A. (1998). Game theory and communication. In Dugatkin, L. A. and Reeve, H. K., editors, *Game theory and animal Behavior*, pages 64–93. Oxford, New York.
- Johnstone, R. A. and Grafen, A. (1993). Dishonesty and the handicap principle. *Animal Behaviour*, 46:759–764.
- Johnstone, R. A. and Norris, K. (1993). Badges of status and the cost of aggression. *Behav. Ecol. Sociobiol.*, 32:127–134.
- Just, W. and Morris, M. R. (2003). The Napoleon complex: why smaller males pick fights. *Evolutionary Ecology*, 17:509–522.
- Kim, Y.-G. (1995). Status signalling games in animal contests. *J. theor. Biol.*, 176:221–231.

- Krebs, J. R. and Dawkins, R. (1984). Animal signals: mind-reading and manipulation. In *Behavioural Ecology: an evolutionary approach*, pages 380–402. Sinauer, second edition.
- Leimar, O. (1997). Repeated games: a state space approach. *J. theor. Biol.*, 184:471–498.
- Maynard Smith, J. (1982). *Evolution and the theory of games*. Cambridge University Press, Cambridge.
- Maynard Smith, J. (1988). Can a mixed strategy be stable in a finite population? *Journal of Theoretical Biology*, 130:247–251.
- Maynard Smith, J. (1991). Honest signalling: the Sir Philip Sydney game. *Animal Behaviour*, 42:1034–1035.
- Maynard Smith, J. (1994). Must reliable signals always be costly? *Animal behaviour*, 47:1115–1120.
- Maynard Smith, J. and Harper, D. (2003). *Animal Signals*. Oxford University Press, New York.
- Maynard Smith, J. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246:15–18.
- McNamara, J. M., Barta, Z., and Houston, A. I. (2004). Variation in behaviour promotes cooperation in the prisoner’s dilemma game. *Nature*, 428:745–548.
- Mitchell, M. (1998). *An introduction to genetic algorithms*. MIT Press, Cambridge, Massachusetts, 1st paperback edition edition.
- Nash, Jr., J. F. (1951). Non-cooperative games. *Annals of Mathematics*, 54:286–295.
- Nowak, M. (1990). An evolutionarily stable strategy may be inaccessible. *Journal of Theoretical Biology*, 142:237–241.
- Orzack, S. H. and Hines, W. G. S. (2005). The evolution of strategy variation: will an ess evolve? *Evolution*, 59:1183–1193.
- Owens, I. P. F. and Hartley, I. R. (1991). “trojan sparrows”: evolutionary consequences of dishonest invasion. *Am. Nat.*, 138:1187–1205.
- Parker, G. A. (1984). Evolutionary stable strategies. In Krebs, J. R. and Davies, N. B., editors, *Behavioural Ecology: an evolutionary approach*, pages 30–61. Blackwell, London, 2 edition.

- Reichert, S. (1998). *Game theory and animal behaviour.*, chapter Game theory and animal contests., pages 64–93. Oxford University Press, New York.
- Searcy, W. A. and Nowicki, S. (2005). *The evolution of animal communication: Reliability and deception in signaling systems.* Monographs in Behavior and Ecology. Princeton University Press, Princeton, NJ.
- Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4(1):25–55.
- Selten, R. (1980). A note on evolutionarily stable strategies in asymmetric animal contests. *J. theor. Biol.*, 84:93–101.
- Selten, R. (1983). Evolutionary stability in extensive two-person games. *Math. Soc. Sci.*, 5:269–363.
- Selten, R. (1988). Evolutionary stability in extensive two-person games — correction and further development. *Math. Soc. Sci.*, 16:223–266.
- Sumida, B. H., Houston, A. I., McNamara, J. M., and Hamilton, W. D. (1990). Genetic algorithms and evolution. *Journal of Theoretical Biology*, 147:59–84.
- Számádó, S. (2000). Cheating as a mixed strategy in a simple model of aggressive communication. *Anim. Behav.*, 59:221–230.
- Számádó, S. (2003). Threat displays are not handicaps. *J. theor. Biol.*, 221:327–348.
- Takada, T. and Kigami, J. (1991). The dynamic attainability of ess in evolutionary games. *Journal of Mathematical Biology*, 29:513–529.
- Taylor, P. D. and Jonker, L. (1978). Evolutionary stable strategies and game dynamics. *Math. Biosci.*, 40:145–156.
- Thomas, B. (1984). Evolutionary stability: states and strategies. *Theor. Pop. Biol.*, 26:49–67.
- Thomas, B. (1985a). Evolutionary stable sets in mixed-strategist models. *Theor. Pop. Biol.*, 28:332–341.
- Thomas, B. (1985b). On evolutionary stable sets. *J. Math. Biol.*, 22:105–115.
- Tinbergen, N. (1952). Derived activities; their causation, biological significance, origin and emancipation during evolution. *Quarterly Review Biology*, 27:1–32.

Zahavi, A. (1975). Mate selection — a selection for a handicap. *J. theor. Biol.*, 53:205–214.

Zahavi, A. (1977). The cost of honesty (further remarks on the handicap principle). *J. theor. Biol.*, 67:603–605.

Appendix 2.A Stability concepts

The following are definitions of the solution concepts used in this paper.

Nash equilibrium A Nash equilibrium (Nash 1951) is a set of strategies, one for each player, such that no player can increase their payoff by unilaterally switching to another strategy. In a symmetrical game such as E85, if a player choosing strategy J in a population where all other players play strategy I receives a payoff of $E(J, I)$, then strategy I is a Nash equilibrium if:

$$E(I, I) \geq E(J, I) \quad \forall J \neq I \quad (5)$$

The Nash equilibrium definition allows for the possibility that strategy J is a neutral alternative to I . It may score equally well, but not better. A Nash equilibrium is presumed to be stable even if J scores equally well, on the assumption that players do not play J due to the application of rational foresight. Strict Nash equilibria are stable without such biologically implausible foresight.

Strict Nash equilibrium At a strict Nash equilibrium (Harsanyi 1973) no player can unilaterally switch to another strategy without decreasing their payoff:

$$E(I, I) > E(J, I) \quad \forall J \neq I \quad (6)$$

An ESS is a less restrictive subset of Nash equilibria than strict Nash, but still requires no rational foresight to maintain stability.

Evolutionary stable strategy Maynard Smith (1982) defined an ESS as follows:

either

$$E(I, I) > E(J, I) \quad \forall I \neq J \quad (7)$$

or

$$E(I, I) = E(J, I) \quad \text{and} \quad (8)$$

$$E(I, J) > E(J, J) \quad \forall I \neq J \quad (9)$$

Condition (7) is that I is a strict Nash or, failing that, an ESS may be a Nash equilibrium (8) that meets the additional condition (9) that the strategy is a better alternative against a potential neutral invader than that neutral invader is against itself. This “second condition” means that a strategy may invade by drift, but that it must be selected against by the original strategy if the invader becomes appreciably common in the population.

ES Sets relax the “second condition” (9) and allow for equilibria composed of a set of mutually inter-invading strategies whose proportions are free to

drift.

Evolutionary stable sets A set of Nash equilibrium strategies L is an ES Set (Thomas 1985b) if:

$$E(I, J) > E(J, J) \quad \forall I \in L, J \notin L \text{ and} \quad (10)$$

$$E(I, J) = E(J, J) \quad \forall I, J \in L \quad (11)$$

ES Sets vs. Mixed equilibria An evolutionarily stable set is only superficially similar to the concept of a mixed Nash equilibrium or a mixed ESS. Mixed equilibria, either Nash or ESS, are compositions of two or more pure strategies at precise proportions. Any deviation of the population away from this optimal mixture, either in terms of population composition in the case of a polymorphic population of pure strategists, or a uniform population of probabilistic mixers, will be returned to the optimal mixture in the case of an ESS, or lead to non-equilibrium strategies in the case of a mixed Nash. Either way, these mixed equilibria form a single equilibrium point in the strategy space. On the other hand, an ES Set produces a line, area, or volume in the strategy space which contains an infinite number of different pure strategy mixtures which are all co-equilibria.

Appendix 2.B Performance of the GA on other tasks

To test the ability of the genetic algorithm to find the ESS to a simpler game we ran a minimally-modified version on several versions of the Hawk-Dove game (Maynard Smith 1982) without role asymmetry (Hurd 2006). This game has a simple mixed ESS when $V < C$, to play Hawk with probability $p = V/C$, and a pure Hawk ESS when $V > C$. We ran three versions, in which V and C were varied so that the ESS was to play Hawk with 16, 83 or 100% probability.

We ran ten replications for each of the three variable combinations, the GA parameters were kept as close as possible to those used in the E85 simulations: population size of 100, simulation length of 500 generations, 5 encounters per individual per generation, and mutation rate of 0.005. Initial seed of ESS players was not used (all initial populations were randomly generated). No crossover operator was used (since the chromosome consisted of a single locus (Hawk or Dove allele) none could be implemented anyway). The percentage of the population playing Hawk was averaged over the last 200 generations of each run.

When the ESS was to play Hawk with 16% the GA resulted in a mean of $17.4 \pm 0.6\%$, when the ESS was 83% the GA produced a mean of $80.0 \pm 0.7\%$, when the ESS was 100% the GA produced a mean of $97.9 \pm 0.3\%$. The bias towards 50% evident in these results can be explained by the fact that mutation will increase the percentage of the least common strategy. Indeed, when we did a further set of 10 simulations for the 83% Hawk ESS with mutation set to 0, the mean outcome was 82.9% Hawk.

We also ran a number of simulations of the E85 GA in which we held the opponent strategy constant at the ESS. From a 0% ESS seed starting point, the population evolved to the ESS in about 100 generations. These results show that the genetic algorithm is capable of finding ESSes when they exist, have large basins of attraction and other non-ESS optima do not exist.

Appendix 2.C Sensitivity of the results to methods and parameters

To test for sensitivity of the major results to the simulation parameter values we ran a set of simulations in which we varied the following model parameters: length of simulation (i.e. number of generations), population size, number of fights per individual per generation, and the use of a crossover operator to make reproduction sexual. Each manipulation was simulated 50 times at each of two levels of initial ESS seed, 0% and 25%. Mutation rate was held constant at 0.002 (a value favorable to the ESS), and all other parameters were held constant at the values used in the main simulations.

These simulations are compared to two similar sets with no manipulations (one at each of 0% and 25% ESS seed) with respect to the number of ESS, ES Set, or Other outcomes obtained. The parameter manipulations comprised of: 6 sets of simulations in which we increased the number of fights per individual per generation from 5 to 25, 50, and 100 fights, two sets in which we increased population size to 1000, and four sets which increased the generations simulated from 500 to 1000 and 5000. Finally, we added a uniform crossover operator which formed an offspring chromosome by swapping between parental chromosomes with a fixed percentage (the crossover “weight”) at each locus. For example, if the first parent was AAAAAA and the second was BBFFFF, with a crossover weight of 20%, each locus would be evaluated individually and swapped with a 20% change. Thus, if the second and sixth loci were selected to be swapped, the returned offspring would be ABAAAF and BAFFFA. We ran four sets of simulations, with crossover weights of 20% and 50% (and ESS seed at 0 and 25 %).

None of the manipulations show an improvement in convergence on the ESS (Table 2.3) with the exception of the population size increases at 25% initial seed. Significant improvement in convergence on the ESS was found at both the 1000 ($G(2) = 10.5$, $p = 0.005$) and 2500 ($G(2) = 10.3$, $p = 0.006$) population sizes. The number of Other outcomes did not decrease, but the ES Set outcomes were less likely at these larger population sizes. Note however that no such effect was seen when the population was started from a random point in strategy space. When the ESS seed was 0% the population never evolved to the ESS at these population sizes.

In conclusion, above and beyond the demonstrated analytical stability of the evolutionarily stable set (see Appendix 4), the present results suggest that our discovery of the new solution to the E85 game using the genetic algorithm was not due simply to our specific choice of implementation or parameter values.

Parameter	Value	ESS	ES Set	Other
0% ESS seed				
Baseline	Defaults	1	41	8
# of Fights	25	4	38	8
	50	3	46	1
	100	1	41	8
Population size	1000	0	44	6
	2500	0	49	1
Simulation length	1000	3	41	7
	5000	1	40	9
Crossover	Weight: 20%	0	46	4
	Weight: 50%	1	46	3
25% ESS seed				
Baseline	Defaults	40	7	3
# Fights	25	28	19	3
	50	38	11	1
	100	36	11	3
Population size	1000	45	0	5
	2500	46	0	4
Simulation length	1000	34	14	2
	5000	39	9	2
Crossover	Weight: 20%	28	15	7
	Weight: 50%	23	24	3

Table C.4: Effect of varying genetic algorithm parameters on the simulation outcomes. The parameters varied (and their default values) were: # Fights (the number of fights used to calculate each individual's fitness each generation, 5), Population size (500), Simulation length (the number of generations simulated, 500), Crossover weight (the likelihood of recombination, 0). The only parameter to increase the probability that the ESS would go to fixation in the population was the population size, but only when the initial population was heavily loaded with ESS players. Increasing the population size did not favor the ESS when starting from a random population.

Appendix 2.D Proof that **AAAA is an ES Set

Following Enquist (1985) and Hurd (1997) we can confirm the stability of the All-Attack ES Set analytically using what we have termed “stability tables”. A stability table for a particular strategy lists the payoffs to a player playing against a population of the target strategy (e.g. the ESS or the ES set) and is logically equivalent to performing a dynamic programming optimization (Houston and McNamara 1999) against the target strategy. From the table, we can determine what the best reply to the target strategy is; if the best reply is the target strategy alone, then it is a strict Nash. Here, we present stability tables for two strategies in the All-Attack set, ABAAAA and AAAAAA. The other strategies of the set, BAAAAA and BBAAAA, are identical by symmetry to the two presented, and so we have omitted their stability tables.

(a) Strong					
Display	Opponent display	Opponent behaviour	Behaviour		
			Attack	Pause-Attack	Flee
A	A	Attack	$\frac{1}{2}V - C_0$ [1]	$\frac{1}{2}V - C_0 - F_P$	$-C_0$
	B	Attack	$V - C_{-1}$ [2]	$V - C_{-1} - F_P$	$-C_{-1}$
B	A	Attack	$\frac{1}{2}V - C_0$ [1]	$\frac{1}{2}V - C_0 - F_P$	$-C_0$
	B	Attack	$V - C_{-1}$ [2]	$V - C_{-1} - F_P$	$-C_{-1}$
(b) Weak					
A	A	Attack	$-C_1$ [3]	$-C_1 - F_P$ [4]	$-C_1$ [5]
	B	Attack	$\frac{1}{2}V - C_0$ [6]	$\frac{1}{2}V - C_0 - F_P$ [7]	$-C_0$
B	A	Attack	$-C_1$ [3]	$-C_1 - F_P$ [4]	$-C_1$ [5]
	B	Attack	$\frac{1}{2}V - C_0$ [6]	$\frac{1}{2}V - C_0 - F_P$ [7]	$-C_0$

Table D.5: Stability table for ABAAAA. Payoffs for all possible outcomes while playing against the communicating all-attack strategy ABAAAA. The variables used as in Hurd (1997): V is the value of the resource, C_x is the cost of an escalated fight against an opponent x strength states higher or lower (1 is higher, -1 is lower, 0 is equal), and F_P is the cost of pausing before attacking an opponent that is simply attacking.

From the tables, it is clear that the following conditions will make ABAAAA stable (numbers in brackets refer to the numbered cells in the tables):

$$\frac{1}{2}V > C_0 \text{ and } V > C_{-1} \text{ [1,2], and } F_P > 0 \text{ [3,4,6,7]}$$

Since cells [3] and [5] are equal, the best reply to ABAAAA is **A[A/F]AA —where * is any allele at the signalling loci. To demonstrate

(a) Strong					
Display	Opponent display	Opponent behaviour	Behaviour		
			Attack	Pause-Attack	Flee
A	A	Attack	$\frac{1}{2}V - C_0$	$\frac{1}{2}V - C_0 - F_P$	$-C_0$
	B	Attack	$V - F_A$	V	$\frac{1}{2}V$
B	A	Attack	$\frac{1}{2}V - C_0$	$\frac{1}{2}V - C_0 - F_P$	$-C_0$
	B	Attack	$V - C_{-1}$	$V - C_{-1} - F_P$	$-C_{-1}$
(b) Weak					
A	A	Attack	$-C_1$	$-C_1 - F_P$	$-C_1$
	B	Attack	$V - F_A$	V	$\frac{1}{2}V$
B	A	Attack	$-C_1$	$-C_1 - F_P$	$-C_1$
	B	Attack	$V - F_A$	V	$\frac{1}{2}V$

Table D.6: Stability table for ABAFAA. Payoffs for all possible outcomes while playing against the communicating strategy ABAFAA. The variables used are as in Table D.5.

that a strategy choosing F at the fourth locus is not a member of the ES set, we present its stability table (Table C.4). The best reply to ABAFAA, one of the potential set members, is not itself but a different strategy AAA[A/F]PP. This strategy is therefore not a member of the ES Set. Similar analysis (not shown) demonstrate that the other three strategies with F at the fourth locus (AAAFAA, BAAFAA, BBAFAA) are not ES set members either.

(a) Strong					
Display	Opponent display	Opponent behaviour	Behaviour		
			Attack	Pause-Attack	Flee
A	A (s)	Attack	$\frac{1}{2}V - C_0$	$\frac{1}{2}V - C_0 - F_P$	$-C_0$
	A (w)	Attack	$V - C_{-1}$	$V - C_{-1} - F_P$	$-C_{-1}$
B	A (s)	Attack	$\frac{1}{2}V - C_0$	$\frac{1}{2}V - C_0 - F_P$	$-C_0$
	A (w)	Attack	$V - C_{-1}$	$V - C_{-1} - F_P$	$-C_{-1}$
(b) Weak					
A	A (s)	Attack	$-C_1$	$-C_1 - F_P$	$-C_1$
	A (w)	Attack	$\frac{1}{2}V - C_0$	$\frac{1}{2}V - C_0 - F_P$	$-C_0$
B	A (s)	Attack	$-C_1$	$-C_1 - F_P$	$-C_1$
	A (w)	Attack	$\frac{1}{2}V - C_0$	$\frac{1}{2}V - C_0 - F_P$	$-C_0$

Table D.7: Stability table for AAAAAA. Payoffs for all possible outcomes while playing against the non-communicating all-attack strategy AAAAAA. The variables used are as in Table D.5.

The situation for AAAAAA is slightly more complicated, as the strategy violates pervasiveness such that the best response to AAAAAA is **AA**. As in the previous case, the two alleles at the signalling loci all achieve the

same payoff. The last two loci, which code for responses to the "B" signal are silent, since the target population never uses the "B" signal. The stability table is constructed with the labels "strong" - "s" - or "weak" - "w" - in place of the signals "A" or "B". Both these rows happen equiprobably following an "A" signal, and so the expected payoff is the mean of their values. The ****AA**** silent mutations produces 36 potential invaders (18 by symmetry). The method for separating the ES set members from the non-members proceeds exactly as in the last case. In the interest of space, the details are omitted.

Chapter 3

The Sir Philip Sidney Game

**When will evolution lead to deceptive
signalling in the Sir Philip Sidney
game?**

Steven Hamblin

Department of Psychology
University of Alberta

3.1 Introduction

Biologists have long wondered whether animals communicate “honestly” when using signals such as threats, begging calls, or courtship displays, or whether they sometimes deceive each other (Dawkins and Krebs 1978; Krebs and Dawkins 1984; Johnstone 1997). Game theoretical modelling has been the principal theoretical method for investigating this question (Maynard Smith and Harper 2003; Hurd and Enquist 2005). Possibly the most frequently cited paper on the evolutionary stability of deceptive signalling is that presented by Johnstone and Grafen (1993). This model is a modified version of the Sir Philip Sidney game (Maynard Smith 1991), in which a donor (the receiver) decides whether to contribute a resource to a beneficiary (the signaller). In Maynard Smith’s version, the resource is Philip Sidney’s water bottle, but the usual biological interpretation is parental provisioning, or some other investment in kin.

In the Johnstone and Grafen version, the beneficiary is either more, or less, needy of the resource, and either a closer, or more distant, relative. The donor does not know the beneficiary’s levels of need or relatedness, only whether or not the beneficiary has signalled a request for the resource be donated. The evolutionary stable strategy (ESS) identified by Johnstone and Grafen has the donor give the resource to signallers who signal for it, while beneficiaries of one relatedness class signal for donation only when needy, but members of the second class always signal for donation regardless of their need. Johnstone and Grafen’s ESS is a semi-separating equilibrium (Fudenberg and Tirole 1991). Beneficiaries may be in one of four relatedness class by need state combinations (closely related and needy, closely related and not needy, distantly related and needy, distantly related and not needy) but must signal with one of two signals (request donation, do not request donation). Figure 3.1 depicts the signal and state combinations present in this game.

A perfectly informative, “separating”, equilibrium would require a minimum of four signals, one per class-by-state combination (Note: see Hurd & Enquist, 2005 for a criticism of the use of the terms “honesty” and “dishonesty” to describe semi-pooling equilibria in this game). Since the signal space is not large enough to be separating, any signalling equilibrium must be semi-separating, since signallers in some class-by-state combinations will have to pool on the same signal. Lachmann and Bergstrom (1998) have investigated pooling equilibria in other forms of the Philip Sidney Game, notably those in which both the donor and recipient have varying need states. or both beneficiary and donor signal. They show that, even for a continuous gradient of beneficiary need, that signal pooling will only ever be between neighbouring categories / contiguous value ranges. This is ambiguous, but arguably not “dishonest”. The same sort of result, discrete signals shared by continuous ranges of subjective resource values exists in the conventional threat display game of resource value (Model II in Enquist 1985;

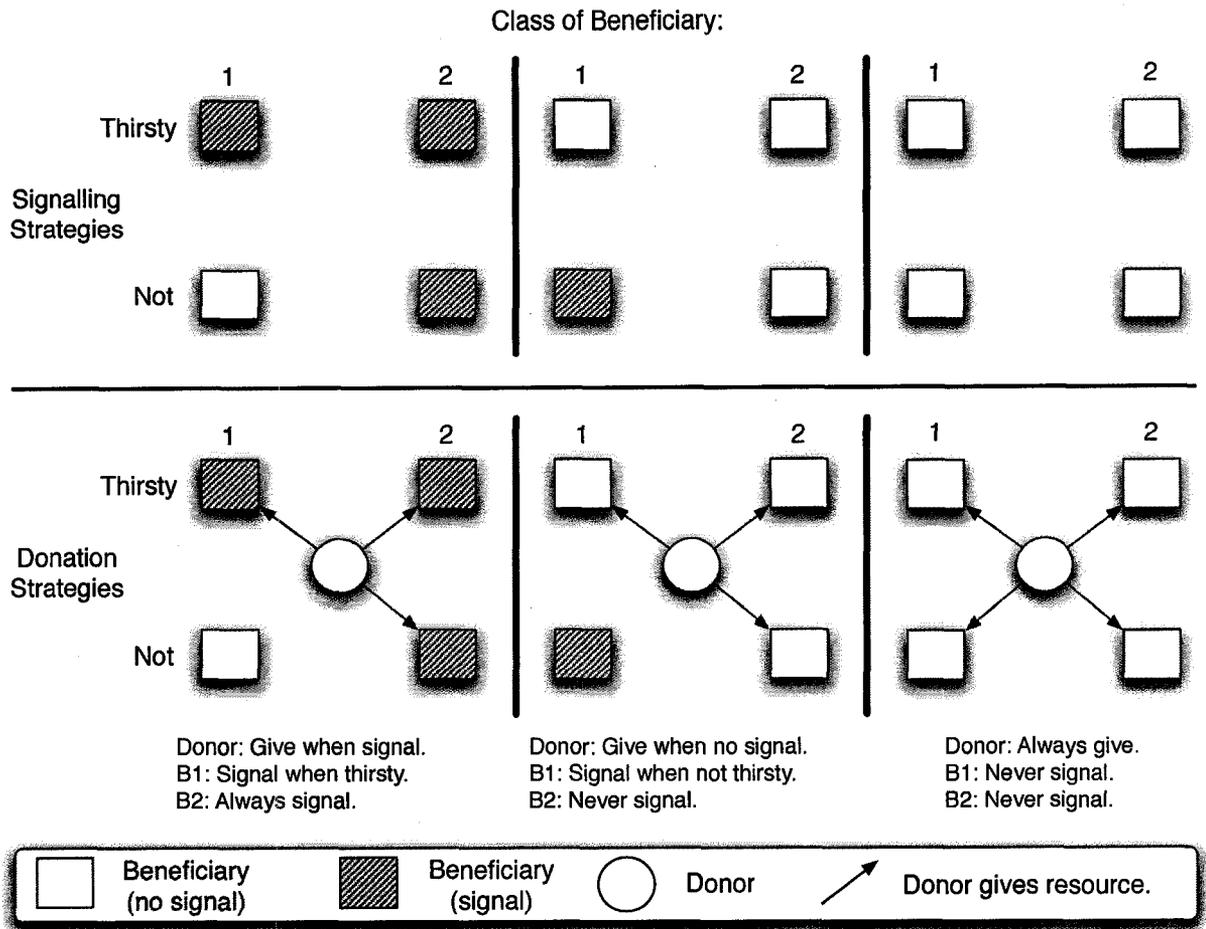


Figure 3.1: Signalling in the Johnstone and Grafen SPS game. Column one depicts the identified ESS to the game, the second column gives an example of another signalling strategy, and the third column is a non-signalling strategy.

Enquist, Ghirlanda & Hurd, 1998). Contrast this with the much more deceptive bluffing threat display model of Adams and Mesterton-Gibbons (1995) in which the highest and lowest RHP signallers pool on one discrete threat display, while signallers of the intermediate RHP range use a different display. What is unique to the Johnstone and Grafen game we investigate here is that, within any play of the game, the signaller may vary in two dimensions simultaneously: relatedness and need.

Within classical game theory, solutions to game theoretic problems typically take the form of Nash equilibria. Under some conditions, Nash equilibria may owe their stability to the application of rational foresight. Such rational foresight has a teleological nature which cannot be reconciled with biological evolution. John Maynard Smith (Maynard Smith and Price 1973; Maynard Smith 1982) defined a refined Nash equilibrium, the ESS, which excludes those cases where rational foresight is required to maintain stability. There is nothing in the definition of an ESS which requires that selection take the ESS to fixation in a population if it is not already very widespread. Therefore, while an ESS is stable once reached, there is no guarantee that populations will evolve towards them (Nowak 1990). An alternative solution methodology, known as evolutionary game theory, analyses exactly this question (Hofbauer & Sigmund 1988; Hammerstein 1998). Positing a reasonable method of strategy replication and population dynamics such as Taylor & Jonker's (1978) replicator dynamics, the space of all possible population compositions can be mapped analytically to show the path along which evolution will take the population of strategies. Since each pure strategy represents an entire dimension of the population's state space, evolutionary game theory analysis is tractable only when there are few pure strategies. Games such as the present form of the Philip Sidney game, where there are four pure strategies for each signaller class and another four for the receiver, result in an unmanageable 64 dimensional space in which to map strategy flow.

We have demonstrated (Hamblin and Hurd 2007) that genetic algorithms may be used to find alternative solutions to signalling games. In this paper, we apply a genetic algorithm to the Johnstone and Grafen Sir Philip Sidney game to test the robustness of the dishonest signalling ESS.

3.2 Methods

3.2.1 Johnstone & Grafen's Sir Philip Sidney game

Johnstone and Grafen's Sir Philip Sidney game has two players, a donor and a beneficiary. The donor has an indivisible resource, which they may choose to give to the beneficiary or retain for their own use. The beneficiary is in the "thirsty" state of greater need with probability p , and "not thirsty" with probability $1 - p$. A move-by-nature assigns the beneficiary to one of two

classes, class 1 or class 2 with probability q and $(1 - q)$. The players are related by a coefficient of relatedness r_i , $0 \leq r_i \leq 1$ with the subscript $i = 1, 2$ indicating which class the beneficiary has been assigned to. Thus, a player's fitness consists not only of their own payoff, but also inclusive fitness benefits of the other player's fitness multiplied by the coefficient of relatedness. If the beneficiary keeps the resource, they ensure survival and receive a payoff of 1, and if they give the resource they receive a payoff of S_B , $0 \leq S_B \leq 1$. If the donor is not thirsty and the beneficiary keeps the resource they receive a payoff of S_D , $0 \leq S_D \leq 1$, while if they are thirsty and the beneficiary keeps the resource the donor dies and receives a payoff of 0. If the beneficiary gives the resource the donor's survival is ensured and the donor receives a payoff of 1. Finally, the beneficiary may choose to give a costly signal. Signalling reduces their fitness by a multiplicative factor of $(1 - t_i)$, $0 \leq t_i \leq 1$ with the subscript i once again indicating the class of the beneficiary. The game is shown in its entire extensive form in Fig. 8 of Hurd & Enquist (2005).

Thus, with two classes of beneficiary there are eight parameters to the model: the two coefficients of relatedness r_1 and r_2 , two signal costs t_1 and t_2 , two survival probabilities S_D and S_B , the probability of being thirsty p and the probability of being a class 1 beneficiary q . To generate payoffs where donors give when they receive a signal, class 1 beneficiaries signal when thirsty and class 2 beneficiaries always signal, Johnstone and Grafen gave parameter values of $r_1 = 0.5, r_2 = 0.2, S_D = S_B = 0.8, t_1 = 0.4$, and $t_2 = 0.3$. However, analysis of the game with these values indicated that their semi-separating equilibrium was not an ESS under these payoffs. Thus, we reformulated the payoffs with $t_2 = 0.1$, which gave the semi-separating equilibrium as an ESS solution to the game.

There are 64 pure strategy profiles in the Johnstone and Grafen SPS game. The donor has four: always give the resource, give the resource only when a signal is received, give the resource only if no signal is received, and never give the resource. Both classes of beneficiary also have four: always signal, signal when thirsty, signal when not thirsty, and never signal. Johnstone and Grafen gave a full formulation of the game in their paper and included payoffs for strategies at the equilibrium, (but did not provide a complete payoff matrix for the game). We provide the reconstructed algebraic payoffs in Table 3.1, while numerical payoffs for the parameter values above can be found in Table 3.2.

	T,S	T,NS	NT,S	NT,NS
Give	$S_D + r_i(1 - t_i), 1 - t_i + S_D r_i$	$S_D + r_i, 1 + S_D r_i$	$S_D + r_i(1 - t_i), 1 - t_i + S_D r_i$	$S_D + r_i, 1 + S_D r_i$
Don't	$1, r_i$	$1, r_i$	$1 + r_i S_B(1 - t_i), r_i + S_B(1 - t_i)$	$1 + r_i S_B, S_B + r_i$

Table 3.1: Algebraic form of the payoff matrix for the SPS game. r_i is the relatedness coefficient for the beneficiary of class i , S_D is the survival chance for the donor, S_B is the survival chance for the beneficiary, t_i is the signalling cost for the beneficiary of class i .

		T,S	T,NS	NT,S	NT,NS
Class 1	Give	1.1, 1	1.3, 1.4	1.1, 1	1.3, 1.4
	Don't	1, 0.5	1, 0.5	1.24, 0.98	1.4, 1.3
Class 2	Give	0.98, 1.06	1, 1.16	0.98, 1.06	1, 1.16
	Don't	1, 0.2	1, 0.2	1.144, 0.92	1.16, 1

Table 3.2: Numerical payoff matrix for the SPS game. T is Thirsty, NT is Not Thirsty, S is Signal, NS is No Signal.

3.2.2 The Genetic Algorithm

Strategies in the genetic algorithm were represented by a "chromosome" consisting of six binary loci. Each of these loci specified the pure local strategy (sensu Selten 1983) for a different information set. All information sets in the game had a corresponding locus, so that the entire chromosome coded for a pure behaviour strategy (sensu Selten 1983). The first two loci represented the donor strategy portion of the chromosome, with the first locus for give / don't give when a signal is received (coded as 0 or 1) and the second locus for give / don't give when a signal was not received. The third and fourth loci similarly represented strategy choices for class 1 beneficiaries, with the third locus for signal / don't signal when thirsty and the fourth locus for signal / don't signal when not thirsty. The fifth and six loci coded the same strategy choices as the third and fourth loci, this time for class 2 beneficiaries. Figure 3.2 graphically depicts the relationship between the game's extensive form and the chromosome.

As an example, the chromosome representing Johnstone & Grafen's semi-separating equilibrium of Give when signal / Signal when thirsty / Always signal under this scheme is $\{1,0,1,0,1,1\}$. The strategy Always give / Signal when not thirsty / Never signal is $\{1,1,0,1,0,0\}$ while Never give / Never signal / Always signal would be $\{0,0,0,0,1,1\}$.

Simulations of the Sir Philip Sidney game with the genetic algorithm were done under conditions similar to those outlined in Hamblin and Hurd (2007): all simulations were done with a population of $n = 100$ for 500 generations, and fitness was calculated in each generation as the mean of 5 plays of the game against other randomly chosen individuals from the population. After fitness was calculated, 5% of the population was killed, with the probability of death inversely proportional to their share of the total fitness of the population. Each death was replaced with a strategy chosen from the top 20% of the population, as ranked by fitness. After reproduction, each locus of every member of the new population was mutated with a probability set by a mutation rate parameter. If a locus was selected to be mutated, the current value was replaced equiprobably with either of the allowable values ($\{0,1\}$) for that locus.

The programming code for the genetic algorithm was based directly on the

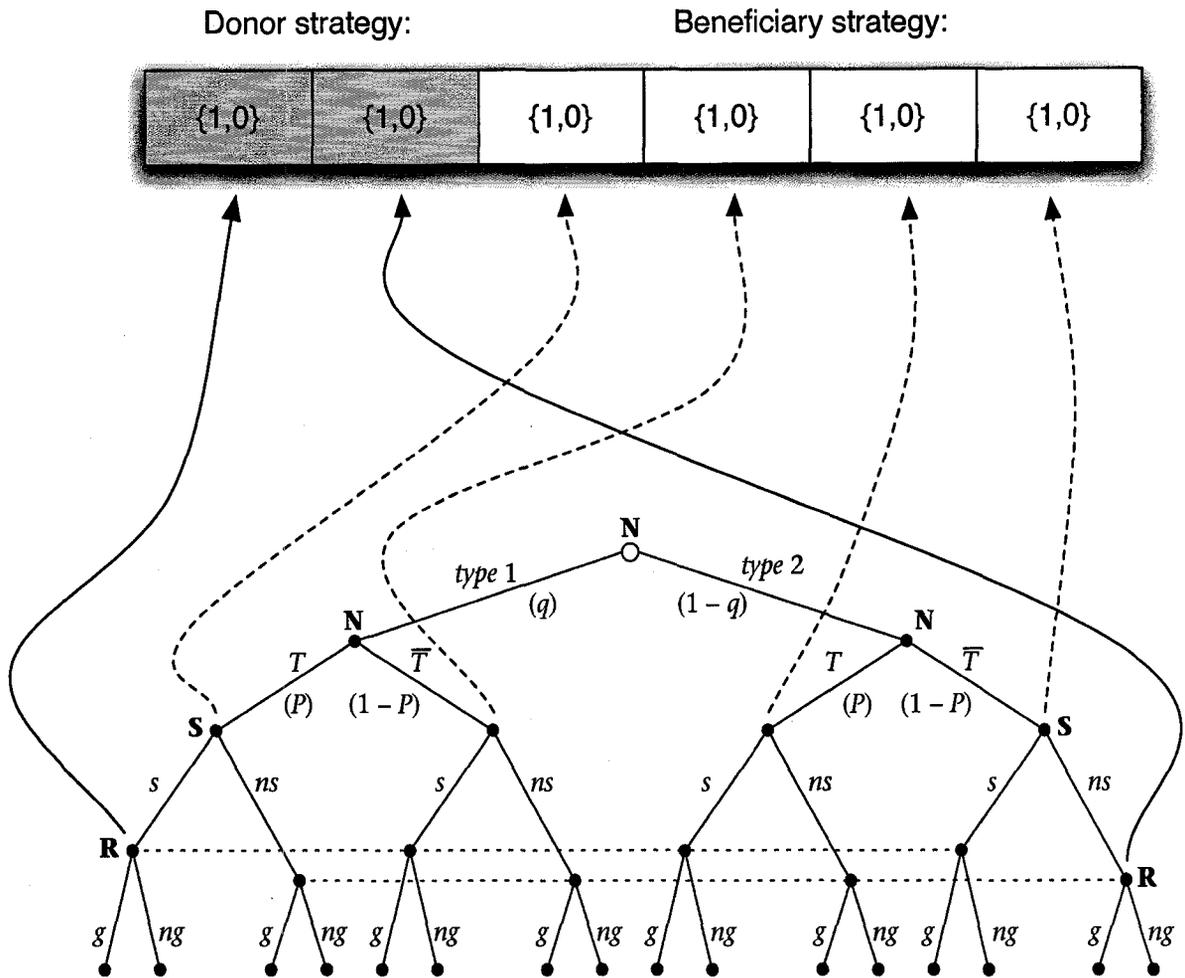


Figure 3.2: Mapping the Sir Philip Sydney game tree to the genetic algorithm's chromosome. Game tree reproduced from Figure 8 of Hurd and Enquist (2005). Note that **N** denotes a move by nature, q is the probability of being class 1 or class 2, P is the probability of being thirsty, **S** denotes the signaller (beneficiary), **R** is the receiver (donor), and T and \bar{T} are thirsty and not thirsty respectively.

code for the genetic algorithm used in Hamblin and Hurd (2007). Extensive testing of the genetic algorithm's performance on simple problems was completed in that paper to provide proof-of-concept of the suitability of this method for finding solutions to game theory models; interested readers are referred to that paper for further details.

To investigate the effects of genetic algorithm parameters on the outcome of this model, we simulated 100 replications at each of several combinations of mutation rate and "seed" (percentage of the population initialized to the semi-separating equilibrium ESS). Mutation rates ranged from 0.001 to 0.01 while seed ranged from 0 to 100% in 5% increments. This results in a total of 21,000 distinct simulations.

3.2.3 Characterizing Outcomes

Each run of the genetic algorithm was classified according to the strategy that was fixated upon, or if no strategy predominated, this was also recorded. A strategy was considered to have gone to fixation if had achieved 80% representation in the population for the last 50 generations (10% of the simulation).

3.2.4 Parameter randomization

To explore the behaviour of the solution set to this game when the parameters of the model were modified, we used Maynard Smith's (1982) stability conditions for ESSs and Thomas's (1985) conditions for the stability of ES Sets under different values of the model's parameters ($r_1, r_2, t_1, t_2, S_D, S_B, P$, and q). We randomized the parameters subject to the constraints which guaranteed the existence of the semi-separating ESS (Johnstone and Grafen 1993):

$$1 - t_1 + r_1 S_D > r_1 \quad (3.1)$$

$$S_B + r_1 > 1 - t_1 + r_1 S_D \quad (3.2)$$

$$1 - t_2 + r_2 S_D > S_B + r_2 \quad (3.3)$$

$$1 + r_1 S_B > S_D + r_1 \quad (3.4)$$

$$qA + (1 - q)B > 0 \quad (3.5)$$

with $A = P[S_D - 1 + r_1(1 - t_1)]$ and $B = S_D - 1 + r_2(1 - t_2)[S_B(1 - P) - 1]$. (Note that Johnstone and Grafen's original parameter values with $t_2 = 0.3$ don't satisfy inequality 3, which is why the semi-separating ESS was not a solution of the game with those values; setting $t_2 = 0.1$ solves this problem and restores the semi-separating ESS as a solution to the game). We performed one million randomizations of the parameters under these constraints by generating sets of eight uniformly distributed random numbers

between 0 and 1 and determining whether they met the conditions given above. The proportion of rejected parameter sets was high: generating a million valid parameter sets required generating over 1.6 billion random combinations, a rejection rate of over 99.9%. Parameter sets that met the constraints were then used to generate the payoff matrix that corresponded to those parameters, and computer analysis listed the solutions to the model under each set of parameters by checking Maynard Smith's and Thomas's conditions for each strategy.

3.3 Results

3.3.1 Genetic algorithm simulations

The outcomes for the simulations are presented in Figure 3.3. Only runs which were performed at mutation rates at or below 0.007 are shown (above this, no appreciable number of strategies reached fixation due to mutation pressure). The ESS went to fixation in nearly 100% of the simulations if the mutation rate was lower than 0.004 and the initial population was comprised of at least 40% ESS players, but unless the semi-separating ESS was seeded into the population, another solution was vastly preferred. This solution, an Evolutionarily Stable Set (ES Set), consists of two strategies: $\{0,1,0,0,0,0\}$, or $\{\text{Give when no signal} / \text{Never Signal} / \text{Never Signal}\}$, and $\{1,1,0,0,0,0\}$ or $\{\text{Always give} / \text{Never Signal} / \text{Never Signal}\}$. This solution was the only other stable solution identified with these parameter values, and was never identified by Johnstone and Grafen in their original publication of this model.

We can see in Figure 3.3 that there is a distinct trend to the results. At low values of seed (0-15%) the ES Set solution dominates the adaptive landscape, and between seeds of 20-40% the ESS's basin of attraction grows until at seeds of 45% or higher it becomes the only stable solution. The results make it clear that unless the population is heavily loaded in favour of the semi-separating ESS, the probability of fixation on the original solution is approximately zero.

3.3.2 Payoff analysis

We can compare the payoffs to the semi-separating ESS at fixation to the payoffs to the ES Set at fixation to determine which is a global optimum. Using Johnstone and Grafen's payoffs the ESS receives the following payoff at fixation:

$$\begin{aligned} \text{Donor: } W &= q[(1 - P)(1 + r_1 S_B) + P[S_D + r_1(1 - t_1)]] + (1 - q)[S_D + r_2(1 - t_2)] \\ \text{Class 1 Beneficiary: } W &= (1 - P)(S_B + r_1) + P(1 - t_1 + r_1 S_D) \\ \text{Class 2 Beneficiary: } W &= (1 - P)(1 - t_2 + r_2 S_D) + P(1 - t_2 + r_2 S_D) \end{aligned} \tag{3.7}$$

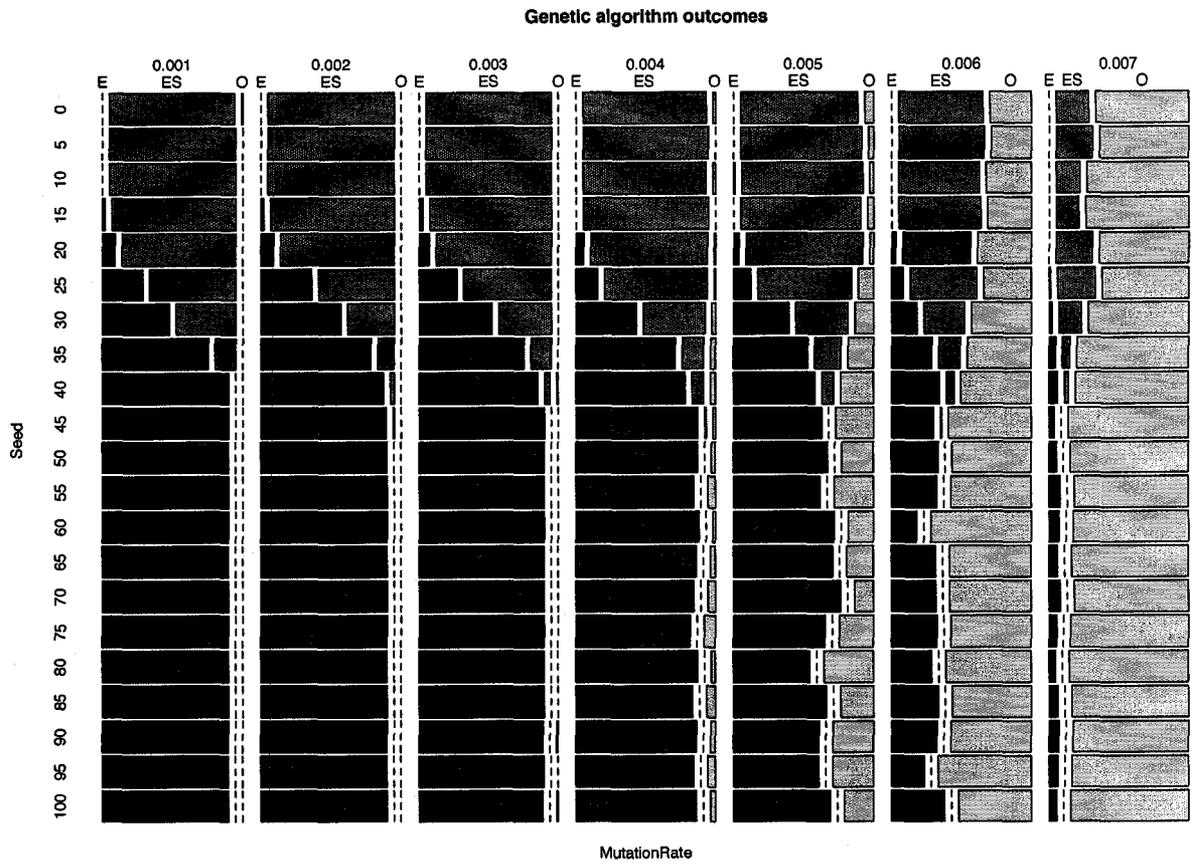


Figure 3.3: Outcomes of the genetic algorithm simulations for various combinations of mutation rate and initial seed. Runs in which the ESS (labeled 'E' on the graph) reached fixation are shown in black, ES Set ('ES') runs are shown in grey, and all other results ('O') are shown in white.

while the ES Set receives:

$$\text{Donor: } W = q[(1 - P)(S_D + r_1) + P(S_D + r_1)] \quad (3.9)$$

$$+ (1 - q)[(1 - P)(S_D + r_2) + P(S_D + r_2)] \quad (3.10)$$

$$= q(S_D + r_1) + (1 - q)(S_D + r_2) \quad (3.11)$$

$$\text{Class 1 Beneficiary: } W = (1 - P)[1 + S_D r_1] + P[1 + S_D r_1] \quad (3.12)$$

$$\text{Class 2 Beneficiary: } W = (1 - P)[1 + S_D r_2] + P[1 + S_D r_2] \quad (3.13)$$

The payoffs to donors who give when no signal is received and donors who always give (the members of the ES Set) are identical at equilibrium since beneficiaries never signal. The payoffs for

$r_1 = 0.5, r_2 = 0.2, S_D = S_B = 0.8, t_1 = 0.4,$ and $t_2 = 0.1,$ are shown in Table 3.3.

	ESS	ES Set
Donor	1.196	1.27
Class 1 Beneficiary	1.12	1.4
Class 2 Beneficiary	1.06	1.16
Total expected payoff	1.155	1.323

Table 3.3: Numerical payoffs to the semi-separating ESS and the ES Set at equilibrium. Total expected payoffs were calculated using the probability q of being class 1 or class 2 beneficiaries and under the assumption that individuals would be donors in half of the encounters and beneficiaries in the other half.

These payoff calculations show that the ES Set is actually a global optimum.

3.3.3 Re-parameterization of the model

Given the constraints used to select the parameters, it was no surprise to find that every randomization contained the semi-separating ESS as a solution to the game. However, we identified a total of five possible regions in the parameter space, each of which contained additional solutions beyond the semi-separating ESS. Table 3.4 lists the solutions for each region. The identified solutions are of two kinds. The first is a variation on the ES Set solution identified by the genetic algorithm, such as the first strategy in region 2 which signals when not thirsty as opposed to never signalling. The second kind is an ES Set (regions 3 and 4) or ESS (region 5) in which donors never give the resource. We attempted to find an intuitive pattern underlying the parameter values for each of the five regions, but were unable to find a simple way to reduce the dimensionality of the parameter set to clarify the relationship between specific parameters and the five regions.

Region	% of outcomes	Chromosome	Strategy - donor	- class 1 beneficiary	- class 2 beneficiary	Solution Type
1	12%	010100	Give when no signal	Signal when not thirsty	Never signal.	ESS
		110000	Always give	Never signal	Never signal.	ESS
		101011	Give when signal	Signal when thirsty	Always signal.	ESS
2	79%	<i>010000</i>	<i>Give when no signal</i>	<i>Never signal</i>	<i>Never signal.</i>	<i>ES Set</i>
		<i>110000</i>	<i>Always give</i>	<i>Never signal</i>	<i>Never signal.</i>	<i>ES Set</i>
		101011	Give when signal	Signal when thirsty	Always signal.	ESS
3	4%	<i>000000</i>	<i>Never give</i>	<i>Never signal</i>	<i>Never signal.</i>	<i>ES Set</i>
		<i>000010</i>	<i>Never give</i>	<i>Never signal</i>	<i>Signal when thirsty.</i>	<i>ES Set</i>
		101011	Give when signal	Signal when thirsty	Always signal.	ESS
4	3%	<i>000000</i>	<i>Never give</i>	<i>Never signal</i>	<i>Never signal.</i>	<i>ES Set</i>
		<i>000010</i>	<i>Never give</i>	<i>Never signal</i>	<i>Signal when thirsty.</i>	<i>ES Set</i>
		010100	Give when no signal	Signal when not thirsty	Never signal.	ESS
		101011	Give when signal	Signal when thirsty	Always signal.	ESS
5	2%	000000	Never give	Never signal	Never signal.	ESS
		101011	Give when signal	Signal when thirsty	Always signal.	ESS

Table 3.4: Regions in the parameter space. Percentage of outcomes refers to the proportion of parameter randomizations which led to the solution set of that region as an outcome. Italicized solutions within a region are members of an ES Set in that region. The semi-separating ESS is a solution in every region, but every region also has an additional ESS or ES Set.

3.4 Discussion

The non-communicating, always-give ES Set was heavily preferred by the genetic algorithm when evolving from a random population. This is in agreement with the results of other models (Rodríguez-Goronés et al. 1996; Lachmann and Bergstrom 1998), suggesting that communication is of dubious value in begging games, and it undermines the original purpose of the model, which was to demonstrate biologically plausible “cheating” in a handicapped signalling situation. Further, random exploration of the parameter space reveals 5 distinct regions in parameter space which all contain the semi-separating ESS along with variations on the first ES Set or variations on a second ES Set solution. This second ES Set solution also undermines the spirit of Sir Philip Sidney games by never giving the resource regardless of the state of the beneficiary. That this strategy set is evolutionarily stable is made even more bizarre by the presence of the original semi-separating ESS in the solution set for these regions of parameter space. Each region contains at least two ESSs (the original semi-separating ESS and a pooling ESS) and one region holds two ESSs and an ES Set which are a mixture of pooling and semi-separating strategies. The dynamics of these other regions in parameter space are likely to be complicated and it is not clear when the “cheating” ESS would be favoured, if ever.

These results demonstrate that even seemingly simple game theoretical models of biological situations may contain dynamics which defy intuition. Sir Philip Sidney games are among the simplest class of signalling models, but the addition of a distinction as simple as the presence of two different classes of beneficiary generates solutions which are contrary to the purpose of the model and result in a sensitivity to starting parameter values that might be termed “chaotic”. Evolutionary dynamics, like those mimicked by genetic algorithms, may prefer these alternative outcomes to the exclusion of the ESS solutions. If models as straightforward as these are subject to such problems, the informal practice of focusing on ESS solutions as the crux of evolutionary stability is severely compromised (see also Hamblin and Hurd 2007). More attention must be paid to evolutionary dynamics, either by simulation or by analytical approaches like evolutionary game theory (Cressman 2003; Hammerstein 1998; Hofbauer and Sigmund 1998; Nowak 2006; Taylor and Jonker 1978) that focus explicitly on population dynamics.

These evolutionary dynamics hint at a larger concern in regards to game theory models in biology. Models are often constructed by conceptualizing a solution to an identified problem and designing a game to produce this solution; in demonstrating that the solution is evolutionarily stable, the identified problem is assumed to have been addressed. This is a sensible method to follow; however, there is no guarantee that any ESSs found will be the only solutions to the game. In simple games, it is easy to specify the entire solution set to the game, but in more complex games with many strategies this can be difficult if not impossible. There is also no guarantee

that these solutions will be preferred as an equilibrium endpoint to the population dynamics; any such arguments must be based on the biological insight into the evolutionary starting points of strategy evolution. Specifying the population dynamics analytically can be done with evolutionary game theory but, as discussed in the introduction, this can be an intractable problem when the number of strategies grows large. As ever, the modeller's most important task is to decide what the essence of the biological problem is before the model is analyzed.

Bibliography

- Adams, E. S. and Mesterton-Gibbons, M. (1995). The cost of threat displays and the stability of deceptive communication. *Journal of Theoretical Biology*, 175:405–421.
- Cressman, R. (2003). *Evolutionary dynamics and extensive form games*. The MIT Press, Cambridge, Massachusetts.
- Dawkins, R. and Krebs, J. R. (1978). Animal signals: information or manipulation? In Krebs, J. R. and Davies, N. B., editors, *Behavioural Ecology: an evolutionary approach*, pages 282–309. Blackwell, Oxford, first edition.
- Fudenberg, D. and Tirole, J. (1991). *Game Theory*. MIT Press, Cambridge, Massachusetts.
- Hamblin, S. and Hurd, P. L. (2007). Evolution's strategies: Genetic algorithms and game theory models. *Animal Behaviour*. in press.
- Hammerstein, P. (1998). What is evolutionary game theory? In Dugatkin, L. A. and Reeve, H. K., editors, *Game theory and animal behaviour.*, pages 3–15. Oxford University Press, New York.
- Hofbauer, J. and Sigmund, K. (1998). *Evolutionary games and population dynamics*. Cambridge University Press, Cambridge, UK.
- Hurd, P. L. and Enquist, M. (2005). A strategic taxonomy of biological communication. *Animal Behaviour*, 70:1155–1170.
- Johnstone, R. A. (1997). The evolution of animal signals. In *Behavioural Ecology: an evolutionary approach*, pages 155–178. Blackwell, Oxford, fourth edition.
- Johnstone, R. A. and Grafen, A. (1993). Dishonesty and the handicap principle. *Animal Behaviour*, 46:759–764.
- Krebs, J. R. and Dawkins, R. (1984). Animal signals: mind-reading and manipulation. In *Behavioural Ecology: an evolutionary approach*, pages 380–402. Sinauer, second edition.
- Lachmann, M. and Bergstrom, C. T. (1998). Signalling among relatives: II. Beyond the Tower of Babel. *Theoretical Population Biology*, 54:146–160.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, New York, NY.
- Maynard Smith, J. (1991). Honest signalling: the Sir Philip Sydney game. *Animal Behaviour*, 42:1034–1035.

- Maynard Smith, J. and Harper, D. (2003). *Animal Signals*. Oxford University Press, New York.
- Maynard Smith, J. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246(1):15–18.
- Nowak, M. (1990). An evolutionarily stable strategy may be inaccessible. *Journal of Theoretical Biology*, 142:237–241.
- Nowak, M. A. (2006). *Evolutionary Dynamics: Exploring the Equations of Life*. Harvard University Press, Cambridge, Massachusetts.
- Rodríguez-Goronés, M. A., Cotton, P. A., and Kacelnik, A. (1996). The evolution of begging: Signaling and sibling competition. *Proceedings of the National Academy of Sciences of the United States of America*, 93(25):14637–14641.
- Selten, R. (1983). Evolutionary stability in extensive two-person games. *Mathematical Social Sciences*, 5:269–363.
- Taylor, P. D. and Jonker, L. B. (1978). Evolutionarily stable strategies and game dynamics. *Mathematical Biosciences*, 40:145–156.
- Thomas, B. (1985). On evolutionarily stable sets. *Journal of Mathematical Biology*, 22:105–115.

Chapter 4

Conclusion

The results I present here demonstrate a method for the modelling of biological problems using game theory. In applying genetic algorithms to these two well known game theory models, I have shown that genetic algorithms can provide solutions for models that are too complex to solve analytically. Furthermore, the use of genetic algorithms has helped to highlight three distinct, but related, problems with the way that game theoretical models are currently used in biology. First, game theory models with an intuitive ESS solution may be obscuring population dynamics which make the ESS is unreachable. Second, as game complexity increases, the number of potential solutions to the model increases as well, and these alternative solutions may be preferred by the population dynamics. This possibility has been largely ignored by biologists who have focused on creating models which contain the desired solution without fully exploring the resulting dynamics of the model. Finally, game theory models currently used in the study of animal communication (and in behavioural biology in general) are impossibly unrealistic from a biological standpoint. Any effort to increase their biological realism will inevitably lead to analytically intractable models, but if we are to use game theory to adequately model biological questions, an increase in realism is necessary. It is the argument of this thesis that the use of genetic algorithms provide a reasonable solution to these problems.

My results suggest that the population dynamics of many game theory models will be unlikely to ever evolve to the ESS (Nowak 1990). These dynamics come from alternative solutions which naturally arise when the complexity of these games rises above that of the simplest models, such as Hawk-Dove (Maynard Smith and Price 1973; Maynard Smith 1982). In games such as those in chapters 2 and 3, the authors have solved the model with the sole aim of finding an ESS which supports the theoretical goal of the model. This practice is common (examples include Grafen 1990; Maynard Smith 1991; Adams and Mesterton-Gibbons 1995; Cushing 1995; Kim 1995; Maynard Smith and Harper 1998; Számadó 2000; 2003; Dubois et al. 2004), and it has the virtue of showing whether a particular strategy can resist invasion once established in a population. Yet if the authors of game theory models wish to claim that their model solutions represent strategies which are attracting as well as evolutionarily stable, then I suggest that this way of solving models must be re-evaluated in favor of a procedure which shows that the population dynamics will justify the dynamic attractiveness of a solution. Genetic algorithms provide an easily-implemented way of providing evidence of this, should the model be too complex to map the dynamics analytically.

The results in Chapter 2 highlight the second problem in biological game theory models, an issue that I explore further in Chapter 3. As game complexity grows, the formal conditions for an ESS become harder and harder to meet, and the potential for multiple solutions to the game grows. Games of anything resembling plausible complexity will either not have an ESS, or will have multiple solutions (e.g. Cressman 1992; 2003). It was this latter problem which the genetic algorithm brought to my attention when I

analyzed the E85 model: the genetic algorithm discovered an entirely new solution which was not published in the original analysis of the game, a solution which had a much greater attractive power in the state space. In other words, populations which did not start off composed entirely (or mostly) of ESS players would not evolve to the ESS but would instead preferentially evolve to the alternative solution. This solution is an Evolutionarily Stable Set (Thomas 1985; Cressman 1992; Balkenborg and Schlag 2001; Cressman 2003), which can be thought of as a set of strategies which would be ESSs if all the other members of the set were removed. This ES Set challenges the original intention of the model, as it is an entirely non-communicating strategy; the original ESS of the E85 game was meant to demonstrate stable conventional signalling, and my work demonstrated that although conventional signalling is stable in this game, it will not be preferred by the population dynamics.

The eventual goal of mathematical models must be to obtain as realistic a depiction of the essence of a problem as is possible while still maintaining an appropriate level of abstraction. Current game theory models in biology are still too simplistic and unrealistic. As Maynard Smith noted twenty-five years ago Maynard Smith (1982), too much attention has been paid to modifications of models like Hawk-Dove when such games contain strategy sets (i.e. sets of possible phenotypes) which are hopelessly naïve. Communication games, such as those analyzed in this thesis, exemplify the direction that game theory models must take to approach a reasonable level of realism. To begin, there are some features of communication that must be added if we are to produce models which model communication correctly. All current models depict each player producing a single signal in succession, or all players signalling simultaneously, and this must be broadened to include the opportunity for multiple signalling opportunities, simultaneously or successively. In addition, simultaneous moves by all players is a nearly universal assumption in biological games; this is a product of the matrix formulation of most games, and it is also completely unrealistic. Aside from those things which I feel must be included to allow communication games to depict real communication situations, there are a host of other features which communications games have included at some time. Among these are:

- Signals which may or may not affect payoffs. These may be handicapped (Grafen 1990) versus conventional (Enquist 1985) signals, cheap-talk models (e.g. Silk et al. 2000), and so on.
- Strength states, often referred to as RHP or Resource Holding Potential (Hurd 2006).
- Role asymmetries (Maynard Smith and Parker 1976; Hammerstein 1981). For example, a strategy choice in a communication game might require specifying the choice of action when the player is the "owner" of

a territory or an "intruder". These are often *correlated* with payoffs (i.e. the role that a player is placed in has an effect on the payoff received).

- Uncorrelated asymmetries. These are asymmetries which are not relevant to the issue modelled by the game and have no effect on the payoffs, but which can be used to settle a conflict (Maynard Smith and Parker 1976). Examples might include owner-intruder (if the asymmetry has no effect on payoffs, i.e. if the owner has no advantage because of ownership), which animal arrived first, the position of each animal with respect to the sun, and so on.

This is by no means an exhaustive list, but adding in even a small subset of the necessary and possible features to a game theory model can lead to a combinatorial explosion in the number of possible strategies. For example, by adding another state variable or signal choice to the E85 game (Chapter 2), we can easily end up with over ten million strategies. A game with a strategy set of this size is unlikely to be easily analyzed with pencil and paper, but it is likely that we must face up to games of this complexity in order to create game theory models which properly capture the realities of the problem being modelled.

When solving games of reasonable complexity is too difficult to do analytically, we must turn to another method, and I suggest that the results presented in this thesis provide further evidence in support of our previous call for the use of genetic algorithms to solve game theory models (Hamblin and Hurd 2007). Future research in this area is required to further prove the applicability of genetic algorithms to game theory models of all levels of complexity. For example, the genetic algorithm approach is as yet untested on game theory models with continuous strategy spaces (both the E85 and Sir Philip Sydney games analyzed here have discrete strategy spaces). The genetic algorithm approach also holds promise for modeling extensions to traditional game theory models, such as the effect of population structure and the dynamics of sub-populations. This work establishes the suitability of genetic algorithms as a solution technique for complex game theory models, and next it will be time to apply them to biological models of greater complexity.

Bibliography

- Adams, E. S. and Mesterton-Gibbons, M. (1995). The cost of threat displays and the stability of deceptive communication. *Journal of Theoretical Biology*, 175:405–421.
- Balkenborg, D. and Schlag, K. H. (2001). Evolutionarily stable sets. *International Journal of Game Theory*, 29:571–595.
- Cressman, R. (1992). Evolutionarily stable sets in symmetric extensive two-person games. *Mathematical Biosciences*, 108:179–201.
- Cressman, R. (2003). *Evolutionary dynamics and extensive form games*. The MIT Press, Cambridge, Massachusetts.
- Cushing, B. S. (1995). When a hawk can damage a dove - an extension of game theory. *Journal of Theoretical Biology*, 175:173–176.
- Dubois, F., Giraldeau, L.-A., Hamilton, I. M., Grant, J. W. A., and Lefebvre, L. (2004). Distraction sneakers decrease the expected level of aggression within groups: A game-theoretic model. *The American Naturalist*, 164(2):E32–E45.
- Enquist, M. (1985). Communication during aggressive interactions with particular reference to variation in choice of behaviour. *Animal Behaviour*, 33:1152–1161.
- Grafen, A. (1990). Biological signals as handicaps. *Journal of Theoretical Biology*, 144:516–546.
- Hamblin, S. and Hurd, P. L. (2007). Evolution's strategies: Genetic algorithms and game theory models. *Animal Behaviour*. in press.
- Hammerstein, P. (1981). The role of asymmetries in animal contests. *Animal Behaviour*, 29:193–205.
- Hurd, P. L. (2006). Resource holding potential, subjective resource value, and game theoretical models of aggressiveness signalling. *Journal of Theoretical Biology*, 241:639–648.
- Kim, Y.-G. (1995). Status signaling games in animal conflicts. *Journal of Theoretical Biology*, 176:221–231.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press, New York, NY.
- Maynard Smith, J. (1991). Honest signalling: the Sir Philip Sydney game. *Animal Behaviour*, 42:1034–1035.

- Maynard Smith, J. and Harper, D. G. C. (1998). The evolution of aggression: Can selection generate variability? [and discussion]. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, 319:557–570.
- Maynard Smith, J. and Parker, G. A. (1976). The logic of asymmetric conflicts. *Animal Behavior*, 24:159–175.
- Maynard Smith, J. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246(1):15–18.
- Nowak, M. (1990). An evolutionarily stable strategy may be inaccessible. *Journal of Theoretical Biology*, 142:237–241.
- Silk, J. B., Kaldor, E., and Boyd, R. (2000). Cheap talk when interests conflict. *Animal Behaviour*, 59:423–432.
- Számádó, S. (2000). Cheating as a mixed strategy in a simple model of aggressive communication. *Animal Behaviour*, 59:221–230.
- Számádó, S. (2003). Threat displays are not handicaps. *Journal of Theoretical Biology*, 221:327–348.
- Thomas, B. (1985). On evolutionarily stable sets. *Journal of Mathematical Biology*, 22:105–115.