

University of Alberta

The Viability of Skyrms' Divide the Cake Model

by

B. Vance Schamehorn



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

Department of Philosophy

Edmonton, Alberta

Spring, 2006



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: 0-494-13753-3

Our file *Notre référence*

ISBN: 0-494-13753-3

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

Dedication

I dedicate this work to my wife Iride and son Stephen, who have stood by me through what turned out to be a much longer process than I ever expected, and encouraged, supported, and inspired me for every minute of it. I could never have finished this work without them.

Abstract

In *Evolution of the Social Contract* (Skyrms, 1996), Brian Skyrms uses an evolutionary game-theoretic model to explore the social convention of equal division or fair sharing. Observing the global predominance of the fair division equilibrium in a certain version of this model, Skyrms suggests this provides a starting point for an evolutionary explanation of the actual predominance of this social convention.

This model has been criticized by other authors and various extensions or modifications proposed, many of which fail to exhibit the global predominance of the fair-division equilibrium, thus putting the applicability of Skyrms' argument into question.

This document examines Skyrms' model, several arguments against it, and some of the proposed modifications. The model is subsequently extended even further in order to discover whether accounting for other considerations or placing it in an alternate framework (particularly an extended evolutionary haystack-model) will allow the global fair division equilibrium to resurface.

Acknowledgements

I would like to express my sincere gratitude to my thesis advisor Wes Cooper of the Philosophy Department at the University of Alberta for his excellent guidance and encouragement for so many years now. In addition, my thanks goes to my other thesis advisor, David Sharp, for first introducing me to this area of research and for first encouraging me to enter graduate studies.

Also, my thanks go to Alex Rueger, John Simpson, and Adam Morton for reviewing draft versions of this document or pointing out new articles to check and new avenues to explore. I'd also like to extend my thanks to the Philosophy department at the University of Alberta for providing one of the most friendly and supportive academic environments one could hope to find.

Finally, I'd like to acknowledge my uncle, Lyle Schamehorn, for his support and encouragement over these last few years.

Table of Contents

Introduction	1
Chapter I - Evolutionary Ethics	
1.1 The Is/Ought Problem and the Naturalistic Fallacy	4
1.2 The Universality and Possibility Problems	6
Chapter II - Evolutionary Stable Strategies and Skyrms' Model	
2.1 Evolutionary Stable Strategies	8
2.2 ESS and Fisher's Gender Ratio Argument	9
2.3 Skyrms' Model	10
2.4 Comparison of Skyrms' Model and Fisher's Argument	15
Chapter III - Criticisms of Skyrms' Model	
3.1 D'Arms, Batterman, and Gorny	17
3.2 Barrett, Eells, Fitelson, Sober	18
3.3 Replicator Dynamics	19
Chapter IV - Coding Project and Modeling	
4.1 Description of the Coding Project	23
4.2 Full Description of Models and Results	24
4.2.1 Skyrms' Model	28
4.2.2 DBG-Style Model	31
4.2.3 New Cost Calculations for DBG-style Model	33
4.2.3.1 Linear Cost Growth	36
4.2.3.2 Greedy Cost Advantage	37
4.2.3.3 Results for New Cost Calculation	38
4.2.4 BEFS-Style Model	40
4.2.4.1 BEFS-Style Model with Cost Calculations	43
4.3 Synopsis of Chapter IV	45
Chapter V - Correlation	
5.1 Re-examining Correlation	47
5.2 Evolution of Correlation	48

5.2.1 Results with DBG-Style Model	50
5.2.2 Results with BEFS-Style Model	55
Chapter VI - Multi-level (Group) Selection	
6.1 Group Selection and Altruism	60
6.2 Wright's Trickle Model	63
6.3 Haystack Model	65
Chapter VII - Coding Project Part II	
7.1 Description of the Group Selection Coding Project	68
7.2 Trickle Model	69
7.3 Haystack Model	73
7.4 Synopsis of Chapter VII	79
Chapter VIII - Conclusions and Findings	
8.1 Synopsis and Conclusions	84
Footnotes	87
Bibliography	90
Appendix 1 - Details of Figures	92

List of Figures

Figure Title(s) ¹	Page
2.1 Payoffs for Players in Divide-the-cake	11
2.2 Results of Uncorrelated Divide-the-Cake Evolution	12
2.3 Results of Correlated Divide-the-Cake Evolution with Correlation Factor 0.2	14
4.1 Results of Uncorrelated Divide-the-Cake	26
4.2 Averaged Results of Uncorrelated Divide-the-Cake	27
4.3 Averaged Results of Skyrms' Model Correlation 0.0	30
4.4 Averaged Results of Skyrms' Model Correlation 0.2	30
4.5 DBG-Style Model with Coefficients <0.0, 0.2, -0.2>	32, 35
4.6 DBG-Style Model with Coefficients <0.0, 0.2, -0.2> with Cost Factor 0.3	35, 39
4.7 DBG-Style Model with Coefficients <0.0, 0.2, -0.2> Per-Attempt Cost	39
4.8 BEFS-Style Model <0.0, -0.2, 0.2>	41
4.9 BEFS-Style Model <0.0, -0.5, 0.5>	41
4.10 BEFS-Style Model <0.0, -0.75, 0.75>	41
4.11 BEFS-Style Model with Revised DBG-Style Cost (Factor 0.3) <0.0, -0.2, 0.2>.....	43
4.12 BEFS-Style Model with Revised DBG-Style Cost (Factor 0.3) <0.0, -0.75, 0.75>	43
4.13 BEFS-Style Model with Per-Attempt Cost (Factor 0.0444) <0.0, -0.2, 0.2>	44
4.14 BEFS-Style Model with Per-Attempt Cost (Factor 0.0444) <0.0, -0.75, 0.75>	44
4.15 BEFS-Style Model with Per-Attempt Cost (Factor 0.0444) <0.0, -0.95, 0.95>	44
5.1 Target Strategies (M=Modest, F=Fair, G=Greedy)	50
5.2 Results for DBG-Style Model with Random Individual Correlation ...	51, 52

5.3	Average Final Correlation for DBG-Style Model with Random Initial Individual Correlation Greedy (Cutoff 0.01)	52
5.4	Average Final Correlation for DBG-Style Model with Random Initial Individual Correlation Fair (Cutoff 0.04)	52
5.5	Results for BEFS-Style Model with Random Individual Correlation ...	55
5.6	Average Final Correlation for BEFS-Style Model with Random Initial Individual Correlation Fair (Cutoff 0.02)	56
5.7	Average Final Correlation for BEFS-Style Model with Random Initial Individual Correlation Greedy (Cutoff 0.0001)	56
5.8	Average Final Correlation for BEFS-Style Model with Random Initial Individual Correlation Greedy (Cutoff 0.01)	56
6.1	Wright's Trickle Model	63
6.2	Maynard Smith's Haystack Model	66
7.1	Trickle Model using Skyrms' Populations (Uncorrelated) Rate 0.01 ...	71
7.2	Trickle Model using Skyrms' Populations (Uncorrelated) Rate 0.02 ...	71
7.3	Trickle Model using Skyrms' Populations (Uncorrelated) Rate 0.03 ...	71
7.4	Haystack Model using Skyrms' Populations (Uncorrelated) Dispersal 10, $\delta = 0.5$	76
7.5	Haystack Model using Skyrms' Populations (Uncorrelated) Dispersal 15, $\delta = 0.5$	76
7.6	Haystack Model using Skyrms' Populations (Uncorrelated) Dispersal 25, $\delta = 0.4$	76
7.7	Uncorrelated Skyrms' Model	81
7.8	Region for which Prob(all-fair) > Proportion(initial fair)	81
7.9	Region in which Prob(all-fair) > Proportion(initial fair) for Skyrms' Model (Correlation 0.04)	82

Introduction

Since Charles Darwin's *Origin of Species* was published 1859, evolution has become a pre-eminent framework for explaining many facets of animal behaviour and physiology. Evolutionary theory has also been widely used as the lens under which human behaviours have been scrutinized, including moral, social, and ethical behaviours. This has given rise to attempts to bridge the gap between the natural sciences and ethics, using evolution as a means to explain why certain moral and ethical behaviours and dispositions are found. These types of ethical explanations are grouped together under the title 'evolutionary ethics.'

Evolutionary explanations in general are not without their drawbacks, and those of evolutionary ethics are no exception. For instance, evolutionary explanations are often good at providing reasons for why certain characteristics are prevalent, but less successful at explaining *how* those characteristics came to exist in the first place.

One type of evolutionary explanation can help to overcome some of the drawbacks traditionally associated with evolutionary explanations, however. This type of explanation draws on the concept of an evolutionary stable strategy (Maynard Smith, & Price, 1973), and involves demonstrating both how certain strategies can become the endpoint for a wide variety of starting conditions, and how those strategies can then remain in existence for long periods of time. Evolutionary stable strategies are most widely used in evolutionary game theory, which provides a framework in which to model evolutionary interactions over time.

Brian Skyrms proposes the possible beginnings of one such explanation in *Evolution of the Social Contract* (Skyrms, 1996) in the context of fair (equal) division of some limited resource vs. unequal division. His explanation is based on a model using a version of the divide-the-cake game, extended over many generations. Skyrms' most suggestive and powerful version of this model involves the addition of correlated encounters, resulting in an evolutionary stable

strategy with nearly global dominance. This dominant evolutionary stable strategy corresponds to the fair division strategy.

Details of this model have come under some degree of criticism by a number of authors (D'Arms, Batterman, & Gorny, 1998; Barrett, Eells, Fitelson, & Sober, 1999; Tennant, 1999). Most particularly, Skyrms' assumptions concerning the nature and amount of correlation have been questioned, and alternate models proposed which attempt to make the model more realistic. These alternate models generally destroy the global dominance of the fair division strategy, allowing an unequal division strategy to persist in a substantial number of cases. This reduces the effectiveness of Skyrms' argument quite substantially, by weakening the result Skyrms relies on from being nearly universal, to instead being merely somewhat more likely than the alternative.

This document will examine Skyrms' model in detail, and try to determine whether, even given some of the assumptions underlying the alternate models, the domain of the fair-division equilibrium can be once again expanded by further extending the model. Also, it will explore the question of how different evolutionary paradigms (especially group selection) might affect the model. My goal throughout this thesis is limited primarily to strictly examining the model itself, and I avoid where possible discussing Skyrms' larger project of which the model is a part. Although I would have enjoyed exploring some of the ethical considerations touched on in greater detail, size considerations prevent me from covering these in the detail they deserve. Because of some of the strengths inherent in Skyrms' style of argument, it will be helpful to settle some of the controversy surrounding the details of the model itself, in order to determine whether further exploration in this direction is likely to be fruitful.

The first chapter of this thesis will briefly look at evolutionary ethics in general terms, and explore a couple of the problems associated with evolutionary ethics. The second chapter will introduce evolutionary stable states, describe Skyrms' version of divide-the-cake, and explore how it answers some of the problems discussed in the first chapter. The third chapter will then detail some of

the criticisms of this model, and introduce the alternate models mentioned earlier. The fourth chapter will describe a large-scale computer programming project designed for use with this thesis, a new method for presenting the results, and then demonstrate this new method by showing the results of some of the alternate models from the second chapter. The fifth chapter discusses correlation in the model, and explores an extension of some models in which correlation is allowed to vary by individual. The sixth chapter introduces and discusses the level of selection issue in evolution and group selection. It continues to describe certain evolutionary frameworks (especially the haystack model) illustrating how differing levels of selection can change the outcome. Chapter VII applies some of the frameworks discussed in the previous chapter to Skyrms' model and examines the results. Finally, Chapter VIII summarizes and discusses the findings and some of their limitations.

This thesis proceeds differently than a typical thesis in philosophy. I began this topic and method of exploration largely because I wanted to be surprised by the results. At the start, it was in no way clear whether the objections to Skyrms' model would persist to weaken it under different assumptions or frameworks, or even what kinds of changes to the model might prove interesting or useful. So, I settled for beginning by simply having a plan for proceeding, but no idea what the end result would be. As a result, this thesis is the culmination of several very different lines of exploration, including a few false starts, a reinvented wheel or two, and dead ends that were looked at and later abandoned. Also, a large amount of the effort put into producing this thesis occurs behind the scenes: a single image might represent the end result of over a thousand hours of computer programming, testing, and simulation time. These things combine to make this thesis seem less precisely focused than some at first glance -- less methodical. I have deliberately attempted to leave some of the exploratory flavour that went into producing this document intact in its presentation, and so I ask the reader to bear with me while I present and tie these somewhat different strands together.

Chapter I - Evolutionary Ethics

1.1 The Is/Ought Problem and the Naturalistic Fallacy

Evolutionary ethics refers to various ways of looking at morality as an outgrowth of the process of natural selection. In these views, at least some moral predispositions are best understood as being instilled in human beings as a result of evolutionary forces, rather than by divine revelation or purely rational means (Schroeder, 2005). This does not necessarily entail that moral sense or moral predispositions are somehow actively beneficial within the process of natural selection, by increasing individual fitness for instance.

Evolutionary ethics can be seen as an attempt to bridge the gap between the natural sciences and ethics. However, there are two somewhat distinct branches of ethics, namely normative ethics and descriptive ethics. Normative ethics attempts to find moral principles that can be used to distinguish between ‘good’ actions, which we should do, and ‘bad’ actions, which we should not do. That is, normative ethics involves the search for standards of behaviour that people *ought* to follow. Descriptive ethics, on the other hand, attempts to explain what moral principles actually are followed by people and why. Sometimes this distinction is referred to as “prescriptive” versus “reportive” ethics (Woolcock, 1999).

Two of the greatest challenges that must be faced by any evolutionary argument which attempts to provide a normative theory of ethics are the is/ought problem and the naturalistic fallacy. The is/ought problem (usually attributed to David Hume) refers to the difficulty of taking some statement of fact, and deriving a normative claim from it. Statements concerning what *is* the case, and statements concerning what *ought* to be the case seem quite distinct. Since evolutionary statements are only of the former variety, deriving normative claims becomes difficult since their claims are of the latter kind. Something needs to bridge the is/ought gap if evolutionary ethics is to provide a normative ethics.

The naturalistic fallacy (made famous by G.E. Moore) is related to the

is/ought problem, and refers to the attempt to define the morally 'good' with reference to something in the natural world. Many theories in evolutionary ethics which attempt to be normative define moral goodness (hence that which we *ought* to do) as something like survival of the species or maximizing evolutionary fitness, both of which are naturalistic claims. These definitions of good are used in an attempt to provide a solution to the is/ought problem. For example, if strategy S maximizes evolutionary fitness, and maximizing evolutionary fitness is morally good (that is, we ought to maximize fitness), then we ought to follow strategy S. If successful, this argument would allow an evolutionary argument to provide a basis for a normative claim. However, there is a difficulty in defining the good in such a manner. If goodness is equated with some naturalistic property, then questions concerning the goodness of that property ought to be redundant, somewhat as the question "is a bachelor unmarried?" is inherently void. For example, if what we ought to do is, say, maximize fitness, then the question "is maximizing fitness good?" should be an empty question. These sorts of questions do not appear to be empty however, and so equating that naturalistic property with a normative moral principle remains problematic.

I am of the opinion that evolutionary arguments are not able to overcome the is/ought problem nor escape the naturalistic fallacy. In the sense of 'non-normative' implied by the is/ought problem and the naturalistic fallacy, I will assume throughout this paper that evolutionary ethics is essentially non-normative. The questions to be answered by evolutionary ethics are not primarily those concerning what ought to be perceived as morally good, but rather questions concerning what we actually do take to be morally good, and how it is we came to perceive them as such. This is especially appropriate given the limited focus of this paper, which aims to examine Skyrms' model in detail while limiting the examination of Skyrms' wider project in *Evolution of the Social Contract*.

My discussion of the is/ought problem and the naturalistic fallacy has necessarily been somewhat brief, as a complete discussion of all the relevant issues would simply be outside the scope of this paper. However, the vast

majority of authors on the subject seem to agree with this limitation of evolutionary ethics (Danielson, 1992; Woolcock, 1999; Schroeder, 2005; Farber, 1994; Dobzhansky, 1962). I have included this brief discussion largely to limit the scope of discussion concerning Skyrms' model in later chapters by excluding treatment of normative concerns.

1.2 The Universality and Possibility Problems

Two other problems that even evolutionary ethical theories of the purely descriptive sort need to overcome are the Universality problem and the Possibility problem. The Universality problem concerns the fact that many ethical-like statements apply to all individuals, whereas evolution typically works at a more limited scope. An evolutionary argument can explain why a behaviour or trait might be applied to family members in order to increase the distribution of the behaviour in subsequent generations (kin selection). An evolutionary argument might also explain why a behaviour is applied to members of an individual's group (group selection). However, when a behaviour is applied to every individual universally, regardless of relationship to the one holding the behaviour, it is more difficult to explain this in evolutionary terms when it does not unilaterally increase the fitness of the one displaying it. Given a choice between a behaviour that operates at a limited scope (e.g. family only), and the same behaviour operating universally, it is difficult to explain why evolution might prefer the latter, particularly if the behaviour or trait in question potentially disadvantages the one holding it.

Another difficulty faced by evolutionary arguments is that even when they explain why a behaviour or adaptation exists in a population, it does not explain how that behaviour or adaptation is possible in the first place (Nagel, 1986, p. 78). For instance, an evolutionary argument might explain why creatures with the ability to use language will survive, but generally does not explain how language use is possible in the first place. Even if it can be demonstrated why or how a behaviour is possible, very often there are alternative possible routes evolution

might take, and the route actually chosen might be initially very unlikely (Nozick, 1993, p. 116). Although slightly different, these two criticisms are related, and I'm going to refer to them collectively as the Possibility problem. Evolutionary arguments appear limited in such a way that the reason they provide for seeing some particular behaviour or characteristic in a population is purely chance, and other possibilities were antecedently as likely (or even more likely) to occur. If what we are using the evolutionary argument to explain is an ethical behaviour, and it is so prevalent as to be nearly universal, then there remains an explanatory gap. We need to explain *why* that moral behaviour is so prevalent if in fact it is purely the result of chance, and initially perhaps a very remote possibility.

Chapter II - Evolutionary Stable Strategies and Skyrms' Model

2.1 Evolutionary Stable Strategies

Evolutionary stable strategies are a type of evolutionary strategy that can address the Possibility problem and Universality problem of the first chapter. In informal terms, an evolutionary stable strategy is a strategy that, provided most of the individuals in a population follow it, cannot be successfully invaded by another strategy. Put another way, an evolutionary stable strategy is a strategy that has higher reproductive fitness than any other strategy provided most of the individuals are already following it (Maynard Smith, & Price, 1973, p. 15).

The definition of evolutionary stable strategy is often formalized more precisely in terms of evolutionary stable states, which I will abbreviate as ESS. Under this more precise definition, a population in which all members are playing an evolutionary stable strategy is said to be in an evolutionary stable state. A state is a set (or simplex) of proportions of individuals playing different strategies, and the set of all possible states is called the *state space*. An ESS (loosely) is a state that has a higher fitness than all other states within some nearby vicinity in the state space (Taylor, & Jonker, 1978, pp. 146-7). A system in an ESS tends to quickly return to its original state if it is perturbed to a slightly different state, although major disruptions can force the system out of the nearby vicinity which can result in different final states. The nearby vicinity of an ESS (the set of states which lead to that ESS) is called the *basin of attraction* for that ESS. An evolutionary stable state is a more general concept than an evolutionary stable strategy, since the former includes states in which more than a single strategy is being used simultaneously by different portions of a population (called a polymorphism), but which nonetheless remains stable indefinitely.

Although Taylor and Jonker provide a formal definition of ESS which is mathematically rigorous, I will usually be applying the more informal definition of the previous paragraph instead. The reason for this is that their formal definition

is described in terms that make it appropriate for a deterministic system (involving replicator dynamics, which I will describe later), but I will be discussing and using many systems which include a more random component, and the formal mathematical definition will not always be applicable.

2.2 ESS and Fisher's Gender Ratio Argument

One ESS in particular is described by Skyrms and makes an excellent illustration. This relates to Fisher's argument for the reason the gender ratio in most animals is close to 50/50, despite having extremely varied biology and widely differing means of reproduction. It might appear for many types of animals that producing more females would provide an advantage, particularly in species where small number of males mate with many females while most of the males do not mate at all, yet this is almost never seen in actual populations. Fisher's argument for why this is so roughly runs as follows. Suppose we have a population of males and females with some other split besides 50/50. For illustration, let us say it is 90 females per 10 males. In such a population, any mutant which produces more male offspring would have an advantage, since males are more likely to have their genes make up a larger proportion of the subsequent generation. Even if a substantial proportion of the males never breed, the argument still holds, since those that do are even more valuable in fitness terms. A similar argument (in reverse) would apply to a population with 90 males to 10 females. It is only a population with a 50/50 gender ratio which avoids the possibility of mutants with higher fitness, and hence a population with an even split remains stable, while other populations are driven to a 50/50 ratio.² This also means a population with a 50/50 gender ratio is in an evolutionary stable state.

Fisher's argument is important because it demonstrates how the notion of ESS can be used to overcome the Universality and Possibility problems from Chapter I. The 50/50 state is the one a population will converge to no matter which state it originally started from. Ultimately, all populations which have two genders will be in an even ratio unless other forces are acting on it which prevent

this. The Universality problem was that moral behaviours are often universal in scope rather than limited. However, if we give an evolutionary explanation for the moral behaviour which involves an ESS of Fisher's type (i.e. one which all populations tend to fall into), then we've made one step towards explaining the universality as well. Since all members of all populations would be following the same strategy, no distinction would need to be made concerning the closeness of individuals. The nature of the strategy has made it such that there is now a reason for assuming that strategies which lack any distinctions of closeness would be preferred in practice. Since it can be reasonably inferred that *any* other individual will be following the same strategy, it could make sense to drop the distinction and make a universal response, which in turn might be a first step towards explaining a universal moral behaviour.

Also, an ESS of this kind can respond to the Possibility problem. Fisher's argument describes how any population with two genders will tend to fall into a 50/50 split. The argument depends only on the mechanics of evolution, and doesn't need to make any assumptions concerning the means of reproduction, the method for selecting mates, or how offspring are reared. In this regard, Fisher's argument is very abstracted, making it applicable to almost any population. The Possibility problem concerned the difficulty of an evolutionary argument's explaining the final prevalence of some strategy, when there might in fact have been many alternate strategies that were all originally possible. In Fisher's case, there is in fact only one strategy which is the final outcome, no matter where the initial population might have started out in terms of gender ratio.³ If an evolutionary explanation for some behaviour is given which relies on this type of ESS, then the Possibility problem is avoided.

2.3 Skyrms' Model

Brian Skyrms gives us a model based on a bargaining game known as divide-the-cake. In this game, two players divide a cake by independently (i.e. without negotiation) deciding how much of the cake to each demand. These

demands are then totalled, and if they can both be fulfilled simultaneously (i.e. they add up to less than or equal to the entire cake), then each player receives the proportion they demanded. If the demands add up to more than the entire cake, each player receives nothing. Skyrms argues that the most intuitive strategy is to demand exactly $1/2$ the cake. In many laboratory settings, this is also the most common strategy seen.

To explain why the demand $1/2$ strategy is so intuitive, Skyrms examines a model in which players choose from one of three different pure strategies: demand $1/3$ (modest), demand $1/2$ (fair), and demand $2/3$ (greedy). How well players following a particular strategy generally fare depends greatly on the mix of strategies being played in the population. Demand $1/3$ players will always receive their demand no matter whom they are paired with, and so that strategy always has a payoff of $1/3$. However, demand $2/3$ players will only receive payoff of $2/3$ if they are fortunate enough to be paired with demand $1/3$ players. Lastly, players using the demand $1/2$ strategy will only receive a payoff of $1/2$ if paired either with demand $1/2$ players, or demand $1/3$ players. A payoff matrix for the game is given below.

Payoffs for Players in Divide-the-cake

		Player 2		
		demand 1/3	demand 1/2	demand 2/3
Player 1	demand 1/3	1/3, 1/3	1/3, 1/2	1/3, 2/3
	demand 1/2	1/2, 1/3	1/2, 1/2	0, 0
	demand 2/3	2/3, 1/3	0, 0	0, 0

Figure 2.1

Since Skyrms uses the replicator dynamics (described in section 3.3) for his model, fitness of a strategy is determined probabilistically rather than using actual pairing. Once the fitnesses of each of the three strategies are determined, they are used to produce a subsequent population (the next generation), with strategies having higher fitness making up a higher proportion of the subsequent

population. For example, if demand $1/2$ tends on average to get more cake than the other strategies, it will make up a larger proportion of the next generation than the other two strategies. This process is repeated until the mix of strategies in the population converges towards an equilibrium.

Skyrms represents the results of this process using triangular diagrams. In these diagrams, each vertex of the triangle corresponds to one of the three pure strategies in the game. Every possible initial strategy mix (i.e. state) corresponds to a point in the triangle, where decreasing distance to the appropriate vertex shows an increase in the proportion of the population initially playing that strategy. Hence, each vertex indicates a population which starts out entirely playing just one of the pure strategies. Also, the edge farthest from a strategy's vertex indicates populations in which that strategy is initially entirely lacking. This triangle is really just a representation of the state space for the model. Lines and arrows are used to indicate how states change from one generation to the next. A rough copy of such a diagram is given below, demonstrating Skyrms' results.

Results of Uncorrelated Divide-the-Cake Evolution

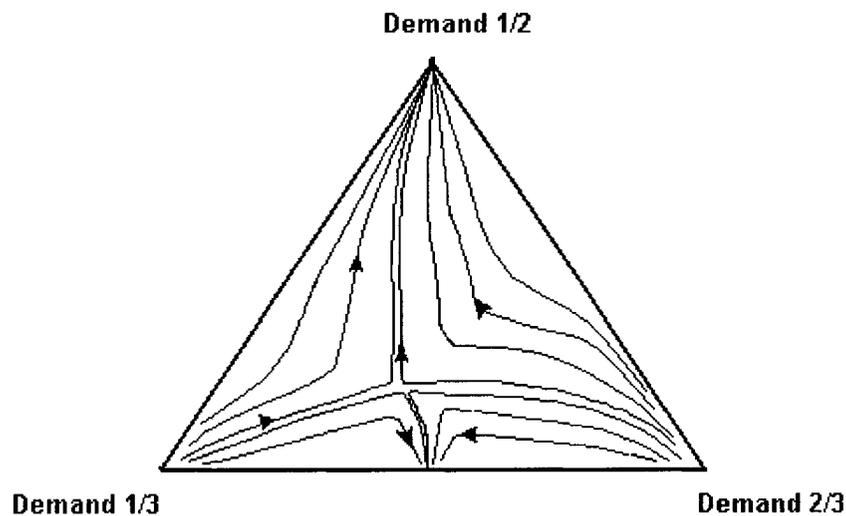


Figure 2.2

The lines and arrows indicate how states change over time. A rough

analogy might be to think of them as indicating the contours of the triangle, so that a ball placed somewhere on it would roll along the line running through that point in the direction the arrow indicated. This is similar to the way one initial state would progressively change over successive generations. The diagram shows a substantial majority of the initial populations move towards the top of the triangle⁴, corresponding to a population with all individuals playing demand 1/2. This is because all the contours in the upper portion of the triangle lead to the top. Our hypothetical ball would roll there and stop once reaching it, since no lines lead away from that point. However, a significant minority of starting states moves towards the bottom centre of the triangle, which corresponds to a polymorphism in which half the population plays demand 1/3 while the other half plays demand 2/3. Each of these two regions is called the basin of attraction for the corresponding equilibrium. Both of these equilibria correspond to an ESS, and thus once entrenched are highly resistant to invasion (perturbation).

Skyrms then goes on to include correlation in the model described above. In this new model, a correlation factor is included which makes a player playing a given strategy more likely to pair with another player playing the same strategy than random pairing would normally allow. This correlation factor changes the outcome of the game quite significantly in many cases. A correlation factor of 0.0 entails purely random encounters, while a factor of 1.0 entails perfect correlation (a strategy is always paired with itself).⁵ A correlation factor as low as 0.2 yields the results shown below in Figure 2.3.

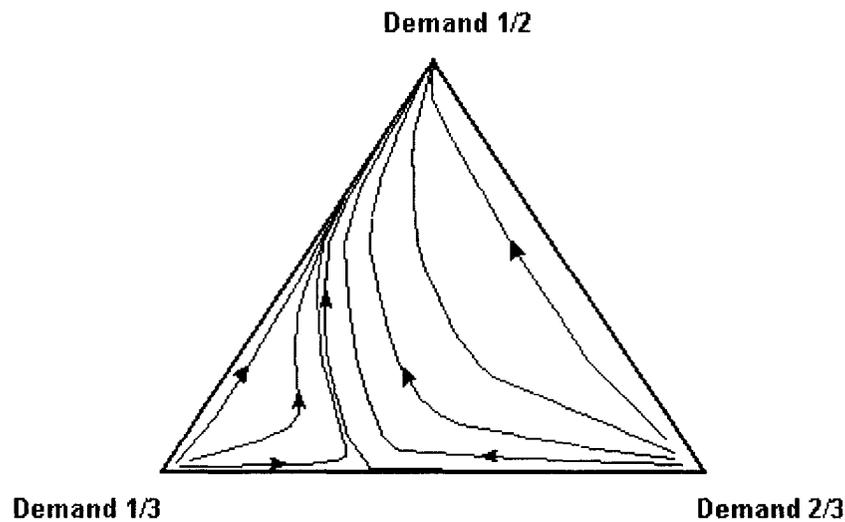


Figure 2.3

The modest/greedy polymorphism has essentially disappeared from the model once a sufficient amount of correlation is included.⁶ The basin of attraction for the all-fair ESS encompasses almost the entire state space, and the all-fair equilibrium is now a global attractor for the game. Skyrms' suggestion is that this model may provide a modest basis for an explanation of the prevalence of fair sharing (equal division) in people. Since the key element in the model is the addition of correlation, how realistic the explanation might be depends very greatly on how likely it is that this correlation actually exists.

There is a long history of using game-theoretic models both in ethics and in evolutionary arguments. Although real situations are considerably more complex and varied than the simpler game-theoretic models, the models nonetheless can provide insight into the real systems. Often the simpler models allow the mechanisms underlying behaviours to be seen clearly which a wealth of extra details would only obscure. Game-theoretic models are also useful because they provide a clear distinction between individuals and situation (Danielson, 1992, p. 30), which can make analysis more straightforward. Although Skyrms' model and the simulations described in this thesis may seem too simple, they can

still demonstrate real behaviour provided they faithfully model all the relevant factors.

2.4 Comparison of Skyrms' Model and Fisher's Argument

Skyrms' model is similar to Fisher's gender split scenario in several key respects. In Fisher's argument, a single equilibrium (ESS) having a 50/50 gender ratio split is the end result of almost every possible initial gender mix. Similarly, Skyrms' model⁷ has a single outcome for all possible initial mixes in which all three strategies are present. Both Skyrms' and Fisher's models have a single global attractor. This was the characteristic of Fisher's argument which gave it some ability to counteract the Universality and Possibility problems of Chapter I, and Skyrms' argument shares this ability.

Equal division (demanding 1/2 the cake) is the most common strategy employed when playing divide-the-cake, regardless of whether the game is being played with family, friends, or strangers. This can remain the case even when the game is changed in such a way that demand 1/2 might be disadvantageous, such as in ultimatum games (Skyrms, 1996, p. 111). The Universality problem questioned why this global strategy application should be seen. However, since the all-fair equilibrium is a global attractor, this would make the assumption that the other player will demand 1/2 highly likely, and hence one's best possible response is to also demand 1/2.

Also, the Possibility problem is avoided, since the all-fair equilibrium is the only antecedently possible outcome in the model. Skyrms' evolutionary argument for the prevalence of fair division explains why we do not find an alternative scenario (particularly the greedy/modest polymorphism). This is particularly useful for an evolutionary argument in ethics, since the behaviour (fair sharing, in this case) is taken out of the realm of behaviours that arose purely by chance, and becomes instead one that is almost guaranteed to develop under the correct circumstances.

One of the most useful parallels between Fisher's argument and Skyrms' is

their generality. Fisher's argument is highly abstracted, and doesn't require any specific details on the mechanism of reproduction, gender selection, genetics, and so forth, in order to still be valid and intuitive. Skyrms' argument is also highly abstracted in this regard. It requires only the existence of different division strategies (including demand 1/2), some mechanism permitting strategies having higher average fitness to make up a higher proportion of the subsequent generation, and a modest degree of correlation. Details concerning how individuals meet, reproduce, and pass on strategies to subsequent generations are all rather irrelevant to the success of this model. This makes it applicable not only to biological evolution, but also conceivably to social or cultural evolution as well, in which imitation and learning are the primary methods by which the behaviour is passed from one generation to the next rather than genetics (Skyrms, 1996, p. 11; D'Arms, Batterman, & Gorny, 1998, p. 88).

This higher level of abstraction is one of Skyrms' argument's strengths. Since it can be applied to many widely varied types of populations, and to both biological and cultural evolution, it increases the plausibility that this scenario will occur. The explanations do not require the existence of some particular biological mechanism in order to function correctly, which makes them more robust than many evolutionary arguments are. This generality and robustness comes at the expense of providing concrete details, however, which would serve to make the argument easier to defend from specific criticisms.

Chapter III - Criticisms of Skyrms' Model

3.1 D'Arms, Batterman, and Gorny

Skyrms' model has been criticized by several authors. Many of these criticisms focus on the addition of correlation to the model, and on the details of how that correlation functions. In particular, an article by D'Arms, Batterman, and Gorny points out that Skyrms uses correlation in such a manner that all strategies are positively correlated. Every strategy is more likely to be paired with itself. However, this really seems advantageous only in the case of fair individuals. Modest individuals get their demanded share (and associated payoff) no matter what kind of individual they are paired with, hence there is no incentive for them to prefer pairing with other modests. Even more importantly, greedy individuals get no payoff whatsoever when paired with another greedy individual. Not only is there no incentive for pairing between two greedies, there is in fact an incentive for greedy individuals to *avoid* one another.

As mentioned before, Skyrms is explicit in giving a generalist explanation, and so deliberately avoids including specific details concerning mechanisms whenever possible in order to increase the robustness. This means that Skyrms' justification for the inclusion of the correlation in the model (and why it acts positively on all strategies) is somewhat vague. D'Arms, Batterman, and Gorny suggest that it in the realm of cultural evolution Skyrms-style correlation might be more plausible (D'Arms, Batterman, & Gorny, 1996, p. 93), but even with this reduction of the domain, there still remain problems justifying positive correlation for all strategies.

D'Arms, Batterman, and Gorny propose a new model, which uses a different correlation scheme. In this scheme, modest individuals are uncorrelated, fair individuals are positively correlated, and greedy individuals are negatively correlated. There are some other changes associated with this model, such as a move away from the replicator dynamics and infinite population sizes Skyrms uses, which they argue increases the realism of their model.

They also argue that one of the advantages to being modest (playing demand 1/3) is that no effort need be wasted on finding an appropriate partner. Fair and greedy individuals have an incentive to be more picky in their selection of partners, resulting in the positive or negative correlation respectively. Modest individuals receive their payoff no matter whom they are paired with, and this lack of selectivity ought to be considered a kind of advantage. A more fine-tuned version of their model is explored, which introduces a cost factor into the simulation, designed to penalize strategies based on the degree of correlation they use.

For brevity, I will refer to this type of correlation (zero for modest, positive for fair, and negative for greedy individuals) as DBG-style correlation. Also, any model which uses it will be referred to as a DBG-style model, regardless of whether it includes the cost factor. Similarly, any model which uses Skyrms' scheme of positive self-correlation for every strategy will be referred to as a Skyrms' model.

3.2 Barrett, Eells, Fitelson, Sober

Even the refinement of the correlation proposed in section 3.1, above, fails to be entirely adequate. Barrett, Eells, Fitelson, and Sober point out that if an individual were to arrange the correlation so as to maximize its own advantage, then the scheme needs to be modified further (Barrett, Eells, Fitelson, & Sober, 1999, p. 240). Just as modest individuals have no incentive to seek out individuals of their own strategy or any other strategy, similarly no other strategy has any incentive for avoiding them. Our use of positive correlation for fair individuals is doing precisely this, however. Fair individuals preferentially seeking out other fair individuals amounts to their avoiding modest and greedy individuals both. Similarly, greedy individuals using negative correlation are avoiding other greedy individuals, which increases the likelihood of their being paired with a fair individual. This outcome is equally bad from the standpoint of the greedy individual. Depending on the mix of strategies present in the

population, negative correlation can be less advantageous for greedy individuals than the equivalent amount of positive correlation would be!

The only lethal result (payoff 0) for a fair individual is when it is paired with a greedy individual. Fair individuals should want to avoid this outcome, and simply avoiding being paired with a greedy individual is a more sure method of accomplishing this than seeking out other fair individuals would be for certain strategy mixes. Similarly, greedy individuals would have an incentive to avoid fair individuals just as much as avoiding other greedy individuals. The best strategy for a greedy individual is thus to seek out modest individuals.

This leads to a new correlation scheme, in which encounters for modest individuals remain uncorrelated, those for fair individuals have *negative* correlation with respect to greedy individuals, and greedy individuals have *positive* correlation with respect to modest individuals. As with DBG-style correlation, I will abbreviate this correlation scheme as BEFS-style correlation, and models using it as BEFS-style models.

3.3 Replicator Dynamics

Brian Skyrms argues that use of the replicator dynamics in evolutionary models provides a good starting point for examining actual systems (Skyrms, 1999, pp. 244-5). Neil Tennant, however, describes how under different kinds of assumptions very different results can be obtained (Tennant, 1999). While I will not be examining the particular model Tennant uses, nor his results, it still appears profitable to look at models which do not rely on the replicator dynamics so strongly.

The populations in Skyrms' model change from one generation to another. If the time between generations is very small, the change in the proportion playing a strategy S_i from one generation to another can be modelled by the following differential equation (Skyrms, 1998, p. 385) :

$$\partial p(S_i) / \partial t = p(S_i) \cdot [U(S_i) - U] / U$$

where $p(S_i)$ is the proportion of the population playing strategy S_i , $U(S_i)$ is the

expected payoff of strategy S_i , and U is the average population fitness. If we let $U(S_i | S_j)$ be the payoff received by playing strategy S_i against strategy S_j , then $U(S_i) = \sum_j p(S_j) \cdot U(S_i | S_j)$ and $U = \sum_i p(S_i) \cdot U(S_i)$. This differential equation is known as the replicator dynamics.

Although it seems complex, it is actually relatively simple to implement in computer simulations. One major assumption inherent in the replicator dynamics is that the proportions of individuals using the given strategies are an exact indicator of the frequency with which the strategies are paired. For instance, if a population is composed of 25% fair individuals and 75% modest individuals, then precisely 25% of modest individuals will be paired with other modests, and the remaining 75% will be paired with fair individuals. Effectively, the replicator dynamics models a situation that usually only occurs in infinitely large populations. As a comparison, the odds of getting tails when flipping a fair coin are said to be 50%, but any finite number of coin flips is very unlikely to produce *exactly* 50% tails. The probability of 50% is an idealization based on an infinite number of flips.

Actual evolution (whether biological or cultural) does not actually take place in infinitely large populations. Skyrms draws a parallel between his own argument and that of Fisher, yet Fisher's argument actually requires a move to finite populations in order to explain the observed scarcity of gender ratio polymorphisms. D'Arms, Batterman, and Gorny also suggest that abandoning the replicator dynamics in favour of finite populations is a move towards increased realism. They further demonstrate that Skyrms' results concerning the relative sizes of the two basins of attraction remain intact even when this change is made, so no loss of explanatory power occurs. Their primary reason for the move is that more complex kinds of correlation and interaction can be effectively modeled in a finite population, but using the replicator dynamics would involve mathematical complexity which could make it impossible to reasonably obtain results. Using the more complex correlation schemes would be highly problematic for this reason.

Also, implementing Skyrms' model using the replicator dynamics can lead to the model's behaving counter-intuitively in many cases. For instance, many populations approach equilibrium without ever actually achieving it. As an example, any population composed of some mix which includes modest individuals has the modest proportion decrease on each generation without ever actually becoming zero. We can say a population approaches the all-fair equilibrium, but we can virtually never say a population is composed only of fair individuals. In an actual (finite) population, this usually does not occur, and overwhelmingly fair populations will almost always have other strategies driven to extinction after a sufficient number of generations.

Besides operating on assumptions only compatible with infinite population size, the replicator dynamics are also completely deterministic. Every initial point in the state space moves to one and only one subsequent state upon the succeeding generation. There is no element of randomness involved. This is also quite unlike actual evolution, which allows for random events and accidents. Even Darwin emphasized that natural selection is usually a noisy process in which luck can play a large role (Binmore, 1994, p. 99). As stated in the discussion of the Possibility problem, evolution in general is such that final possibilities which were highly unlikely initially might still come about from time to time. With the replicator dynamics, however, there are no alternate possibilities: a given starting point proceeds to a single outcome in one and only one way, and all other outcomes have probability of zero. We would like our model to incorporate alternate possible evolutionary paths, otherwise we risk the charge that the reason we always come to the all-fair equilibrium is simply because any alternatives have been rendered impossible by the model. Imposing linear change on the model by using the replicator dynamics reduces the model's ability to counter the Possibility problem.

Although some distinctively new behaviours result from the switch to finite populations, I will still refer to any finite population variation as a "Skyrms' model." This is consistent with the terminology used by D'Arms, Batterman, and

Gorny who also make the switch away from replicator dynamics. The most substantial difference between the various models I discuss is the correlation scheme used, so for the purposes of this thesis a Skyrms' model is any one that uses positive self-correlation for every strategy.

Chapter IV - Coding Project and Modeling

4.1 Description of the Coding Project

In order to test Skyrms' model, the modifications proposed by other authors, plus some ideas of my own, some computer code is required. There is a fair amount of public-domain code available which would seem at first to be potentially adaptable. Unfortunately, few of the best code bases seem to include the capacity to simulate group-level selection theories, which I will be examining extensively. Of those that do, all seem designed in such a way as to make interaction between individuals nearly impossible, due to the data abstraction used.

Evolutionary simulation is most commonly used in genetic algorithms, which are primarily problem-solving algorithms. In the interests of efficiency, these algorithms typically drop the individual, focusing instead only on the 'genes,' which represent possible solutions to the problem to be solved, and which are used to determine fitness. In our case, this approach is unworkable, since the fitness of an individual is also the result of pairing with another individual (i.e. dividing-the-cake), rather than being solely inherent in the genes. That is, the composition of subsequent generations is partly determined by the interactions between individuals, and not just by individual genes (or strategies) per se.

These kinds of individual interactions are more typical of A-Life simulations, and there is plenty of code available for these as well. A-Life (artificial life) encompasses a wide range of computer simulation types. Typically, they involve more agent-based approaches, where interactions between individual agents are simulated in order to study emergent behaviours. I was unable to locate any extendable agent-based A-Life code which included a sufficiently powerful evolutionary component, however. A hybrid approach is required, which allows for a flexible combination of both evolutionary modeling and individual interactions.

Performing very extensive modifications to someone else's code can be a

daunting process at best, particularly when modifying the code to perform in a way the designer never intended. Often, this can be more labour intensive than designing new code, which is what I have elected to do. I have designed and written a very extensive object library in the C++ programming language, called the VE Evolutionary Class Library, which is ideally suited for the kind of hybrid approach I am interested in. It was designed in such a way so as to maximize portability (so that older systems could successfully compile and run simulations), and to be easily adaptable to *both* A-Life and evolutionary simulations.

I had originally intended to append the code and the documentation to this thesis as an appendix, but the code grew to a very large size (well over 140 pages of code plus over 80 pages of documentation) so this plan was abandoned. Instead, the code will be made available for download over the internet, and may be considered as an online component of this thesis. I will be releasing this code into the public domain under the GNU General Public License, which will allow it to be freely copied, modified, or distributed, so long as it is not incorporated into proprietary software. More information on this license may be found in the appropriate text document accompanying the code, or by writing to the Free Software Foundation, Inc., 59 Temple Place - Suite 330, Boston, MA 02111-1307, USA.

4.2 Full Description of Models and Results

I will be using a modeling approach that includes a number of the changes proposed by D'Arms, Batterman, and Gorny. In particular, I will not be using the replicator dynamics. Instead, finite populations will be used for all simulations. Also, the fitness of the different strategies, used to determine the composition of subsequent generations, will not be determined probabilistically (i.e. as a function of the proportions of the different strategies in the current generation). Instead, I will be using the strategy of D'Arms, Batterman, and Gorny once again, and will pair individuals from the population by sampling without replacement, determining fitness at the level of the individual itself. This allows for varying

behaviour of the system,⁸ which in turn allows for a given starting point to lead to a different ESS on different trials. Because of this, multiple runs are performed, and the results averaged, in order to give a clearer indication of the behaviour of a particular starting state.

One other major change I am making is in the way the results are represented. Skyrms used diagrams in which lines and arrows showed the sizes of the basins of attraction as well as the manner in which states changed from one generation to another. However, since the simulations I will be using are designed so that evolution for a given state can vary from one run to another, this is no longer possible. While it is feasible to average multiple runs together and produce a reasonable facsimile, doing so would obscure the fact that along the boundaries, substantial numbers of runs may be evolving in different directions from other runs. The vector method of representation cannot show this uncertainty.

Instead, I will be using colour in order to show the results of the simulations. Each point in the triangular results diagram (corresponding to a given initial state) will be coloured according to the final equilibrium it reaches. I am assigning the primary colour blue to the fair individuals, red to greedy individuals, and green to modest individuals. The proportions of strategies in the final state correspond to the proportions of the three primaries mixed in the colour assigned to the point for that state. For example, the all-fair equilibrium contains only fair individuals, so the colour assigned to any state resulting in the all-fair equilibrium is composed only of blue. Similarly, since the greedy/modest polymorphism is composed of roughly 50% greedy and 50% modest individuals, the colour assigned to a state leading to the polymorphism is coloured a roughly equal mix of both red and green (resulting in a greenish-brown colour). The diagram below shows what this looks like for a Skyrms' model without correlation, and performing only a single run for any given starting point.

Results of Uncorrelated Divide-the-Cake

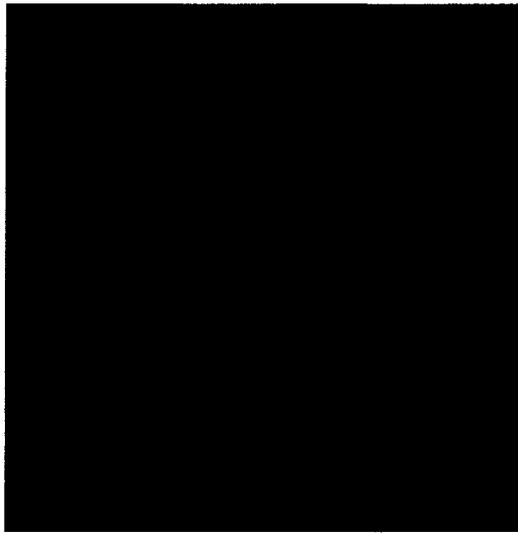


Figure 4.1

The number of starting points used for Figure 4.1 was kept deliberately small (i.e. the granularity was increased) in order to demonstrate how nearby states can randomly end up in very different final states. Some points which finished in the greedy/modest polymorphism can be clearly seen surrounded by points that finished in the all-fair polymorphism. Note that the slight fuzziness surrounding these points (and at the basin edges in general) is due to the interpolation method I use for filling in points for which no evolutionary trial was performed. Later runs will use much finer granularity,⁹ and any blurring introduced due to interpolation will be very small, and typically limited to an area of only one or two pixels. Also note that the corners of the triangle are assigned in the same way that Skyrms does, with the modest vertex at the bottom left, greedy at the bottom right, and fair at the top. I will follow this convention throughout, and omit labelling of the corners on all subsequent triangle diagrams.

There are several benefits to this mode of representation. For instance, it becomes rather simple to visually compare the sizes of the two basins of attraction. Also, shifts in the proportions of strategies in the final equilibrium can be seen as colour shifts in the appropriate basin. Finally, the tendency of points along the boundary of the basins, which I will refer to as the separatrix (Richter,

1986, p. 96), to lead to different equilibria will show up as a mixing of the basin colours along the region of the separatrix when multiple runs are averaged. The following diagram shows the model from Figure 4.1, but with 10 runs averaged together.

Averaged Results of Uncorrelated Divide-the-Cake

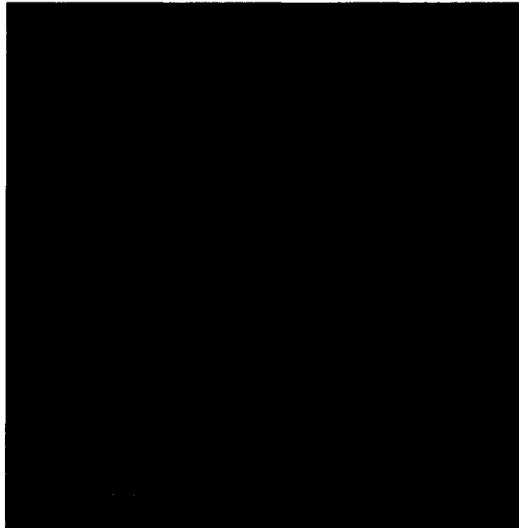


Figure 4.2

As can be seen, the separatrix is much more easily definable once a number of runs are averaged together, but the tendency of some areas to end up in differing basins is also clearly visible. Most of the diagrams in this thesis will use 50 runs averaged together, and a much finer granularity than figures 4.1 and 4.2 use. Also, I have assigned the red, green, and blue primaries to the strategies in such a way that the colours for the two basins have good contrast, and yet still look very similar for persons both with or without colour deficiencies, and I will consistently use this assignment throughout.

Also, the reader may note the very small pure green and pure red areas at the bottom corners of the diagram. This corresponds to the states which initially begin with only modest individuals (in the case of green), or greedy individuals (for red). Since there is no mutation in any of the models used, a population with only modest individuals continues without change indefinitely. The state which

begins with only greedy individuals actually ends up with every individual receiving a payoff of zero, and when the population renormalization procedure of Skyrms' model or the DBG-style model is applied this results in a singularity (dividing strategy payoff of 0 by total payoff of 0). Since none of the authors explicitly mention this, I have elected to renormalize those populations having a total payoff of zero as if a single greedy individual managed to successfully divide the cake, resulting in a new population of all greedy individuals. Later simulations will have other starting states which can die out (besides just overwhelmingly greedy states), so this convention of renormalizing to all-greedy will later be abandoned in favour of colouring these starting states a very dark grey instead.

4.2.1 Skyrms' Model

As described earlier, I am making certain modifications to Skyrms' original divide-the-cake model, such as using finite populations and more realistic pairing. The implementation of Skyrms' model used here is essentially the same method proposed by D'Arms, Batterman, and Gorny. Described simply, individuals are paired randomly, with the correlation value used to increase the likelihood of an individual being paired with another individual using the same strategy. Fitness is calculated after all individuals are paired. Finally, the proportions of the three strategies in the next generation are set according to the proportions of their fitness. In more formal terms, the model uses the following procedure.

- Select one of the individuals in the population at random and remove it from the population.
- Let S_1 , S_2 , and S_3 be the proportion of each of the strategies present in the resulting population (after removal) such that S_1 is the proportion of individuals with the same strategy as the removed individual, and S_2 , and S_3 are the proportions of individuals using the other strategies. $S_1 + S_2 + S_3 = 1$.

- Let e be the correlation factor. Let $S_1' = S_1 + e(S_2+S_3)$ if $e \geq 0$, otherwise $S_1' = (1+e) \cdot S_1$. Let $S_2' = (1 - S_1') \cdot (S_2 / (S_2+S_3))$, and $S_3' = (1 - S_1') \cdot (S_3 / (S_2+S_3))$.
- Choose a random number r such that $r \geq 0.0$ and $r < 1.0$. If $r \leq S_1'$ remove a second individual from the population with the same strategy as the first, otherwise if $S_1' < r \leq S_1'+S_2'$ select a second individual from the population which uses the strategy corresponding to S_2 , otherwise select a second individual from the population which uses the strategy corresponding to S_3 .
- Pair the two individuals removed together in the divide-the-cake game and assign payoff accordingly.
- Repeat the previous 5 steps so long as there are two or more individuals remaining in the population.
- Let M , F , and G be the total payoffs received by all individuals which played modest, fair, and greedy, respectively. Let N be the total size of the original population (before any individuals were removed). Renormalize the population so that $[M / (M+F+G)] \cdot N$ individuals use the modest strategy, $[F / (M+F+G)] \cdot N$ individuals use the fair strategy, and $[G / (M+F+G)] \cdot N$ individuals use the greedy strategy.

I include this more formal description in part to make the differences between this model and later models more explicit and easier to describe. Also, the formal description allows my results to be duplicated by those interested.

As mentioned previously, there are cases in which the final renormalization step fails (i.e. cases where payoffs for every individual is 0). In this event, I will either set G equal to $2/3$, or flag the population as having become extinct for simulations which track population size changes. Note too that since we use finite population sizes, the number of individuals calculated during renormalization must be rounded to whole numbers. This rounding will cause very slightly different results from those of Skyrms in some instances, such as for

initial states with a very large majority of modest individuals. The following two diagrams show my final results for two of Skyrms' models.

Averaged Results of Skyrms' Model

Correlation 0.0

Correlation 0.2

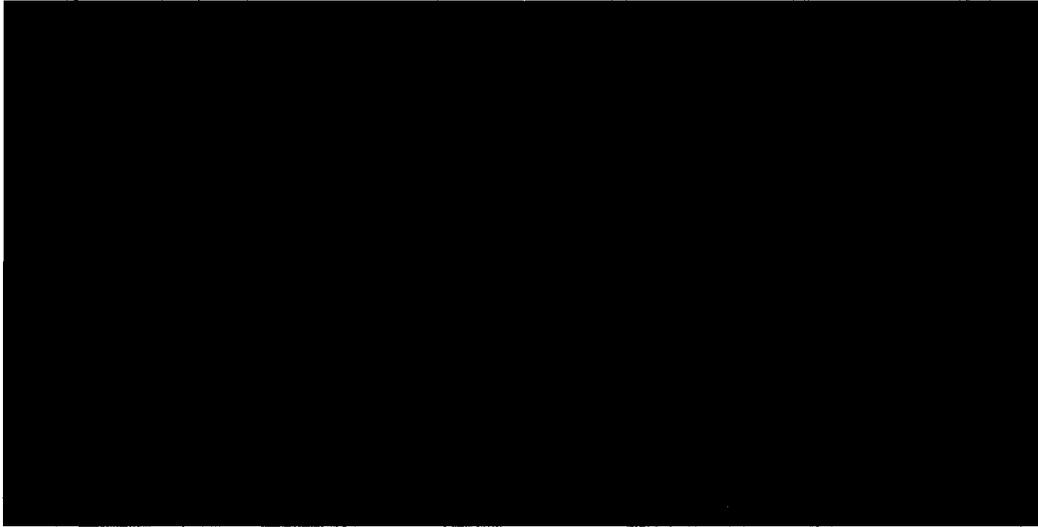


Figure 4.3

Figure 4.4

The second diagram shows practically no basin of attraction for the modest/greedy polymorphism, and the all-fair equilibrium has come to dominate virtually every starting point.

Unlike Skyrms' diagram, a very slight ridge at the bottom still remains, and a careful observation shows a 'speckled' region above it which indicates that some populations do in fact manage to end up in the polymorphism despite being quite some distance from what should be the basin of attraction for it.¹⁰ This is mostly the result of the switch to finite population size and randomness. In the region near the very bottom edge, fair individuals are very uncommon, and there is always some small probability that they will nearly all be paired with greedy individuals rather than other fair individuals or modest individuals. This causes them to decline in proportion very sharply (or become extinct entirely), leaving only modest and greedy individuals. The closer we get to the bottom edge, the more likely this is to occur, especially near the right hand (greedy) vertex.

One way of visualizing this is in terms of Figure 2.3 earlier, which shows the deterministic (replicator dynamics) version of the same model. The arrows in Figure 2.3 represented the change of a population from one state to another over successive generations. In the lower area of the triangle, populations follow closely alongside the bottom edge for a time before moving upwards to the all-fair vertex. When randomness and finite populations are used, the arrows can be thought of as just representing the most *probable* trajectory, but populations can (and do) move in different directions on occasion. Every so often, a population will move downwards toward the bottom edge rather than alongside it, and once there a population is never able to leave that bottom edge. Hence, no matter how strong an attractor the all-fair equilibrium becomes, there will always be some slight fuzziness near the bottom edge of the triangle.

4.2.2 DBG-Style Model

This model is virtually identical to the version of Skyrms' model described in section 4.2.1. The only difference is in how the correlation is handled. Rather than a single correlation coefficient e which acts on all the strategy types, there are now three coefficients e_M , e_F , e_G for the modest, fair, and greedy strategies respectively. The coefficient e_M will always be zero for all DBG-style models, while e_F will be greater than zero and e_G will be less than zero.

The different coefficients come into play in the third step of the model, which was:

- Let e be the correlation factor. Let $S_1' = S_1 + e(S_2+S_3)$ if $e \geq 0$, otherwise $S_1' = (1+e)S_1$. Let $S_2' = (1 - S_1')(S_2 / (S_2+S_3))$, and $S_3' = (1 - S_1')(S_3 / (S_2+S_3))$.

This step needs to be modified as follows in the DBG-style model.

- Let e be the correlation factor corresponding to the strategy of the removed individual. Let $S_1' = S_1 + e(S_2+S_3)$ if $e \geq 0$, otherwise $S_1' = (1+e)S_1$. Let $S_2' = (1 - S_1')(S_2 / (S_2+S_3))$ and $S_3' = (1 - S_1')(S_3 / (S_2+S_3))$.

The change in correlation from a single globally applicable value to three strategy-dependent values makes a substantial change in the behaviour of the model as compared to Skyrms' model. When listing the coefficients used for a DBG-style model (or BEFS-style model later), I will list them as the vector $\langle e_M, e_F, e_G \rangle$. When the coefficients are $\langle 0.0, 0.2, -0.2 \rangle$ the modest/greedy polymorphism continues to maintain a substantial basin of attraction, as shown in Figure 4.5.

DBG-Style Model with Coefficients $\langle 0.0, 0.2, -0.2 \rangle$

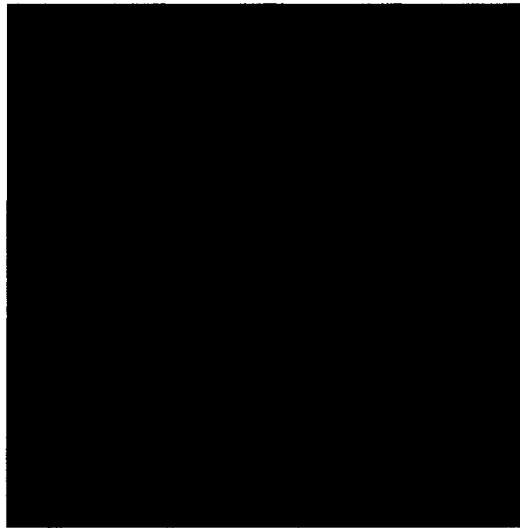


Figure 4.5

D'Arms, Batterman, and Gorny also propose an extension to this model. This extension is described by them as increasing realism (D'Arms, Batterman, & Gorny, 1996, p. 97), and is intended to model the increased costs of an individual's being more selective in choosing a partner to be paired with. In the DBG-style model, fair and greedy individuals are more highly correlated (or anti-correlated) than modest individuals, and it seems reasonable that this would result in longer search times for an appropriate partner under most circumstances. Their modified model adds a cost factor, c , which is used to calculate a reduction of the fitness of each strategy. On each pairing, a cost is calculated such that

cost = $c \cdot e \cdot (1 - S_1)$ if $e \geq 0$, otherwise cost = $-c \cdot e \cdot S_1$. Again, e is the correlation coefficient for the strategy of the individual, and S_1 is the proportion of the strategy of the individual in the population at the time of that pairing. Cost depends on correlation and strategy frequency, and not on the strategy itself, per se. However, modest individuals lack correlation in DBG-style models, and so have the advantage of lower cost because of that. In their modified model, the size of the basin of attraction for the polymorphism becomes even larger, increasing in size roughly 13% when c is 0.3. Brian Skyrms refers to these sorts of cost calculations as a promising area for further investigation (Skyrms, 1999, p. 248).

4.2.3 New Cost Calculations for DBG-style Model

It is unclear from the description provided whether this cost should be applied to only the first individual selected (i.e. the one seeking a partner), or also to the second as well. It is intended to indicate the costs of being more selective in choice of partner, but the article provides no clues as to whether the authors interpret any given pairing as just the first individual selecting the second, or as both individuals selecting one another.

Another difficulty with the cost calculations in the modified DBG-style model is that *negative* payoffs occur in the case of an unsuccessful pairing. This seems rather unintuitive, since the payoffs are added together during the renormalization step to determine the composition of the next generation. Any unsuccessful pairing of a fair or greedy individual can be thought of as reducing the fitness of all other successful fairs or greedies. This changes the nature of the game played, and effectively substitutes a completely new game for the model, rather than remaining true to Skyrms' model. In Skyrms' model, being successfully paired results in a nonzero payoff, while unsuccessful pairing has a zero payoff. I therefore propose a slightly reworked version of the DBG-style cost calculations. This cost is calculated before renormalization, rather than for each pairing. Formally, the final step of the DBG-style model, which was originally

- Let M , F , and G be the total payoffs received by all individuals which played modest, fair, and greedy, respectively. Let N be the total size of the original population (before any individuals were removed). Renormalize the population so that $[M / (M+F+G)] \cdot N$ individuals use the modest strategy, $[F / (M+F+G)] \cdot N$ individuals use the fair strategy, and $[G / (M+F+G)] \cdot N$ individuals use the greedy strategy.

will be changed to now read:

- Let M , F , and G be the total payoffs received by all individuals which played modest, fair, and greedy, respectively. Adjust M , F , and G as follows. Let $M = M \cdot (1 - 3 \cdot c \cdot e_M)$ if $e_M \geq 0$, otherwise $M = M \cdot (1 + 3 \cdot c \cdot e_M)$. Let $F = F \cdot (1 - 2 \cdot c \cdot e_F)$ if $e_F \geq 0$, otherwise $F = F \cdot (1 + 2 \cdot c \cdot e_F)$. Let $G = G \cdot (1 - 1.5 \cdot c \cdot e_G)$ if $e_G \geq 0$, otherwise $G = G \cdot (1 + 1.5 \cdot c \cdot e_G)$. Let N be the size of the original population before any individuals were removed. Renormalize the strategies in the population so that $[M / (M+F+G)] \cdot N$ individuals use modest strategy, $[F / (M+F+G)] \cdot N$ individuals use the fair strategy, and $[G / (M+F+G)] \cdot N$ individuals use the greedy strategy.

I will refer to this as revised DBG-style cost calculations, and it behaves very much like the calculations used by D'Arms, Batterman, and Gorny, except that it is now based on the number of successful pairings. That is, only survivors (with non-zero fitness) have to pay the cost, and nonsurvivors do not modify this cost in any way. The factors 3, 2, and 1.5 above are used to find the number of successful pairings for each strategy from the total payoff. For example, F (total payoff for the fair strategy) equals the number of successfully paired fair individuals multiplied by the fair payoff of $1/2$; multiplying F by 2 simply reverses this and yields the number of successfully paired fair individuals. In simpler terms, the formulas subtract cost-correlation-(number of survivors) from the total payoff for a strategy.

The results with this change to the cost calculations are even more significant than the 13% found with the unrevised version. Figure 4.6 shows the results of the modified cost with a cost factor of 0.3, shown alongside the original non-cost version of the DBG-style model (Figure 4.5) for contrast.

DBG-Style Model with Coefficients <0.0, 0.2, -0.2>

With No Cost Factor

With Cost Factor 0.3



Figure 4.5

Figure 4.6

The motivation for the addition of the DBG-style cost calculations was to add realism to the model by allowing for a slight advantage for modests for being less selective in choosing partners to play with. However, there still remain two problems with both the DBG-style cost calculations and my modified version of cost which reduces the realism of these cost calculations somewhat. The first of these problems is the linear growth of the cost factor with decreasing proportions. The second problem involves the overlooking of one of the primary advantages of the greedy strategy as soon as cost factors are included: greedy individuals can afford to pay more, other things being equal, since they receive a larger payoff from each successful pairing.

4.2.3.1 Linear Cost Growth

A look at the cost calculations shows that they are a linear combination of the cost coefficient, the absolute value of the correlation coefficient, and (for DBG-style cost) the size of the proportion of the population being avoided. This does have the desired effect of increasing the cost depending on how selective a strategy is, and how difficult it is to maintain that selectivity. However, the linearity of the cost seems unrealistic. Consider the case where fair individuals make up 50% of the population. The DBG-style cost using a factor of 0.3 and correlation 0.2 would work out to be 0.03. Now contrast this with the case where fair individuals make up a miniscule 0.0001% of the population. The DBG-style cost in this case is about 0.0599999, which is close to twice as high. Since fair individuals are so scarce in the second case, the cost should realistically be very much higher than it was for the first case. In the first case, the chances of meeting another fair individual after even a few random pairings is quite high, while in the second case it would require around 693,000 random pairings to have even a 50/50 chance of meeting another fair individual. Given this vastly increased difficulty in finding the desired pairing, the cost for the second case should be much higher than merely double what it was for the first case. The DBG-style cost does not realistically portray the increase in difficulty associated with more difficult searches, particularly in cases with very high or perfect correlation (coefficient of 1.0), in which only a specific sort of pairing is permitted.

In order to overcome this problem with linearity, I propose to modify the cost calculations once again. I'll use the original, unmodified DGB-style method of calculating cost for each individual at the time of pairing. However, rather than use a simple linear calculation, I will actually count the number of random pairing attempts made by the first individual, and subtract a fixed amount per attempt. I will refer to this cost calculation method as per-attempt cost.

At some point, the total cost incurred may meet or exceed the potential payoff. How many failures this requires depends on the strategy of the individual. When this occurs, that individual remains unpaired and is assigned a payoff of

zero. Since there is some uncertainty about whether one or both individuals should be paying the cost, I have elected to have only the first individual pay the cost. This seems to me to be the more realistic of the two choices. It seems plausible that there might be some individuals in a population which are found by an appropriate partner before they begin to search for themselves, and such individuals would incur no search costs.

4.2.3.2 Greedy Cost Advantage

In the DBG-style model, each individual chooses a partner strategy type based on a random number r used in the fourth step of the process. This partnering method allows us to correctly model the effects of correlation which can be different for each strategy type. I have mentioned earlier that one of the motivations for introducing cost calculations into the model was to mimic the advantage that modest individuals have by being less selective. In unmodified DBG-style cost calculations, this is accomplished by multiplying the correlation factor in, and since the correlation for modest individuals is zero by default, this makes the cost zero for modest individuals as well.

In per-attempt cost calculations, we have the added complication that modest individuals might be looking for a type of individual that is uncommon in the population at large, and so fail to find a partner with the desired strategy before the costs associated with unsuccessful pairings overtake the payoff for the first individual. This might seem to remove or mitigate the advantage of modest individuals in being less selective. However, this advantage of the modest strategy still persists in per-attempt cost, although to a somewhat lesser degree. Since modest individuals do not use correlation, the strategy type looked for will always reflect only the strategy mix in the current population, so the most common strategy present is always the most common strategy sought out.

One way of looking at original DBG-style cost calculations is that the cost was always zero for a modest individual, no matter how rare the target strategy is in the population at large. Similarly, costs for other strategies were always

nonzero, no matter how slight the correlation or how plentiful the target strategy. In effect, this was like stacking the deck in favour of the modest strategy. This in turn benefited the greedy strategy far more than the fair strategy, since the greedy strategy is heavily reliant on a larger modest proportion. It is not at all surprising that the modest/greedy polymorphism rebounded to such a large extent after the DBG-style cost calculations were introduced.

By blunting this advantage of the modest strategy somewhat, per-attempt cost also allows us to take into account one of the advantages of the greedy strategy, which was overlooked in the original DBG-style cost strategy. Specifically, since greedy individuals get a higher payoff from a pairing, they can afford to pay a higher cost -- being more selective is worthwhile if the rewards for success are correspondingly greater. This serves to increase the realism of cost calculations even more, as well as better fitting the original justification for including cost calculations, which was that the degree of selectivity for a strategy should impact its fitness.

4.2.3.3 Results for New Cost Calculation

Finding an appropriate per-attempt cost to use which makes a reasonable comparison with the original DGB-style cost method is somewhat difficult. The behaviour of the two different cost methods are simply too different for the more extreme strategy mixes (i.e. those near the corners of the triangle). Also, the cost in the DBG-method depends on the correlation being used. I compromised and used a per-attempt cost factor of 0.0444. This allows 7, 11, or 15 pairing attempts for the modest, fair, and greedy strategies, respectively. This is sufficient to give any individual a 99% chance at finding the target partner type in a population composed of an equal mix of all three strategies. The value of 0.0444 also provides (on average) a very similar cost as the DBG-style for both fair and greedy individuals in a population with equal amounts of all strategies and a 0.6666 correlation factor. This correlation may seem high, since we have been using a much lower correlation than that. However, the per-attempt cost increases sharply

as the proportion of the target strategy type decreases, and matching the two strategies at a higher correlation value permits the two methods to approximate one another in a region surrounding the centroid of the triangle, rather than at the centroid only, making the two methods more comparable. Finally, 0.0444 is sufficiently low so that even a modest seeker (which can least afford an extended search), is still more likely than not to be able to pair successfully with a strategy that represents as little as one tenth of the global population. This will avoid accidentally penalizing modest individuals too much.

DBG-Style Model with Coefficients <0.0, 0.2, -0.2>

Modified DBG-Style Cost

Per-Attempt Cost

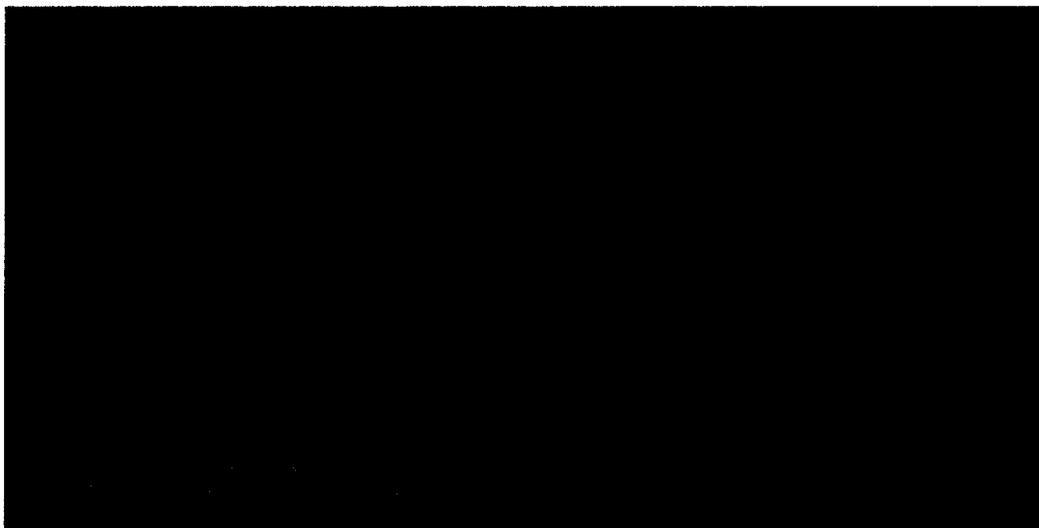


Figure 4.6

Figure 4.7

Figure 4.7 above shows the results I obtained in a simulation using the per-attempt cost (factor 0.0444), alongside Figure 4.6 from earlier for comparison. One thing to note is that the basin of attraction for the per-attempt cost method is substantially smaller. In terms of size, it is slightly smaller even than the case where no cost calculations whatsoever are used (Figure 4.5). This decrease in size occurs because the area that was added to the basin for the modest/greedy polymorphism by DBG-style costs in Figure 4.6 occurs in a region where fair individuals are a minority, perhaps 20% of the population. In this region, per-

attempt cost gives roughly equal chances of successfully pairing to all strategies, whereas DBG-style cost gave a considerable advantage to the modest strategy in this region.

The effects of adding cost calculations to the model are very much like the effects of adding correlations changes: adding more realism or detail can make a substantial difference in how the model behaves overall. It may seem counterintuitive that when cost calculations are redesigned so as to provide some advantage to greedy individuals (longer search times), it is in fact the fair basin of attraction that receives the more substantial benefit. This is partly due to the specific per-attempt cost value chosen, and very small values will produce behaviour much like the non-cost DBG-style model. Very large values on the other hand will rapidly reduce the proportion of modest individuals in most regions of the triangle, and thus rapidly reduce the proportion of greedy individuals which depend on them. In effect, the large increase in size of the basin for the modest/greedy polymorphism that was seen in the original DBG-style cost model was an artifact caused only by its bias towards individuals using the modest strategy. Lessening this bias results in growth of the all-fair basin.

4.2.4 BEFS-Style Model

This model makes some adjustment to the correlation scheme. It uses reasoning similar to that used in the DBG-style correlation scheme, but goes one step further. DBG-style correlation has positive correlation for fair individuals, which results in their avoiding modest individuals as well as greedy (i.e. a higher probability of pairing with another fair entails a lower probability of pairing with a modest/greedy individuals). The BEFS-style model alters this so that fair individuals avoid greedy individuals instead, as this is the only result for the fair strategy which would result in a worse payoff. Similarly, greedy individuals in DBG-style correlation avoid other greedy individuals, when fair individuals would result in just as bad a payout. BEFS-style correlation substitutes for the greedy strategy a positive correlation towards modest individuals, which is the only result

with nonzero payout. This is another attempt to increase the realism of the model by altering the correlation scheme so as to maximize individual advantage.

Barrett, Eells, Fitelson, and Sober suggest that this modification to the correlation scheme should have qualitatively similar results to those seen in the DBG model, resulting in a larger basin of attraction for the modest/greedy polymorphism (Barrett, Eells, Fitelson, & Sober, 1999, p. 240). This turns out not to be the case, and higher correlation values will actually decrease the size of the basin for the polymorphism. Figures 4.8 through 4.10 show the results of the BEFS-style model for increasing correlation values. Note that the correlation values shown still refer to the vector $\langle e_M, e_F, e_G \rangle$ as in the DBG-style model, except the signs have changed because the correlations for the various strategies are now applying to different target strategies.

BEFS-Style Model

$\langle 0.0, -0.2, 0.2 \rangle$

$\langle 0.0, -0.5, 0.5 \rangle$

$\langle 0.0, -0.75, 0.75 \rangle$



Figure 4.8

Figure 4.9

Figure 4.10

Skyrms refers to this decrease of the polymorphic basin as the result of the “overshoot” problem of negative correlation (Skyrms, 1999, p. 249). Consider a population consisting primarily of modest and greedy individuals, in roughly equal numbers. As the correlation of greedy individuals towards modest individuals becomes very high, more and more greedy individuals will be successfully paired with modest individuals. Some modest individuals will still be paired with other modests, however, since they are uncorrelated and select a

pairing strategy based on the current strategy mix. These factors in turn force an increasing number of greedy individuals to be paired with themselves for a mutual payoff of zero, reducing the proportion of the greedy strategy in subsequent generations. The equilibrium thus shifts towards a population where the greedy individuals obtain the same average payout received as that obtained by the modest individuals (namely $1/3$). Skyrms claims this occurs when the population contains twice as many greedy individuals as modest for perfect correlation (i.e. every modest is paired with a greedy). This is actually slightly oversimplified, since the BEFS-style model keeps modest individuals uncorrelated, and so some modests will be paired with other modests. The actual equilibrium would have average payoffs in the same ratio as the proportions, and contains roughly 36.6% modest individuals. Essentially, the equilibrium ‘overshoots’ the ideal proportion mix in which most greedy individuals can be successfully paired.

Although we are not using perfect correlation in the simulation shown above, there is still this shift in equilibrium due to overshoot, which reveals itself as a reddening of the colour of the polymorphic basin in Figure 4.9 and 4.10. The shift of the equilibrium to a population with more greedy individuals than modest reduces the overall efficiency of the modest/greedy polymorphism even further compared to the all-fair equilibrium. This loss of polymorphism efficiency combined with the high efficiency of fair individuals causes the reduction in the size of the basin of attraction for the polymorphism for the high correlation values. In fact, at sufficiently high correlation values, the behaviour of the BEFS-style model is nearly the same as in the original Skyrms-style model with correlation. Since the BEFS-style correlation was argued to be more realistic, this could be construed as providing support for Skyrms’ original argument. However, the correlation coefficients required to eliminate the polymorphism are considerably higher than before, and such high values might prove unlikely in many real-world evolutionary scenarios.

4.2.4.1 BEFS-Style Model with Cost Calculations

Although the basin of attraction shrinks for the modest/greedy polymorphism for the BEFS-style correlation scheme with high correlation values, the addition of cost calculations in the DBG-style model had the opposite effect, increasing the size of the basin. It will be interesting to see how the cost calculations change the results in the BEFS-style model, and to determine whether the polymorphic basin reduction in it can be overcome by the DBG-style cost scheme's support for the modest strategy. The results from two different runs using a DBG cost factor of 0.3 are shown in Figure 4.11 and 4.12 below, the first of which uses the standard correlation values of ± 0.2 , and the second using higher values.

BEFS-Style Model with Revised DBG-Style Cost (Factor 0.3)

$\langle 0.0, -0.2, 0.2 \rangle$

$\langle 0.0, -0.75, 0.75 \rangle$

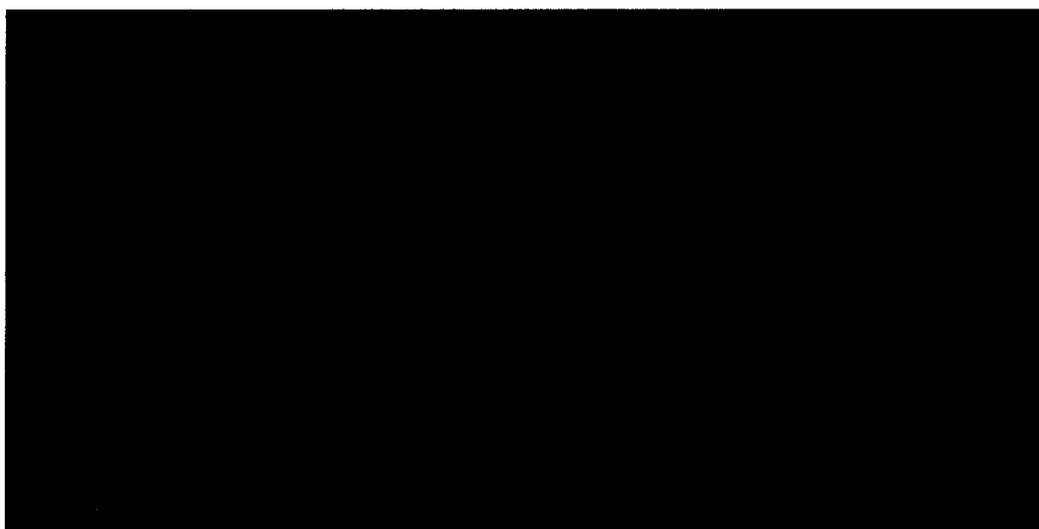


Figure 4.11

Figure 4.12

The modest/greedy polymorphism basin has effectively disappeared in Figure 4.12.¹¹ The reduction of the basin of attraction for the polymorphism as a result of overshoot could not be compensated for with the addition of DBG-style cost calculations. Also, a close inspection shows that the red-shift due to the overshoot equilibrium shift that was seen earlier in Figure 4.9 and 4.10 is absent in Figure 4.12. As discussed in the preceding section, a certain proportion of

greedy individuals were not being paired with modest individuals when correlation was high, and there was a resulting shift in the equilibrium such that greedy individuals outnumber modest individuals. However, the addition of cost calculations to the model has the opposite effect, increasing the ratio of modest individuals to greedy. In Figure 4.12, these two opposing tendencies have nearly cancelled one another out, resulting once again in an equilibrium with a nearly 50/50 split. This shift of equilibrium for the modest/greedy polymorphism towards the modest corner as a result of DBG-style cost calculations also occurred earlier in the DBG-style model, and can be seen when contrasting Figure 4.6 with Figure 4.5 (although the colour difference is subtler).

BEFS-Style Model with Per-Attempt Cost (Factor 0.0444)

$\langle 0.0, -0.2, 0.2 \rangle$

$\langle 0.0, -0.75, 0.75 \rangle$

$\langle 0.0, -0.95, 0.95 \rangle$

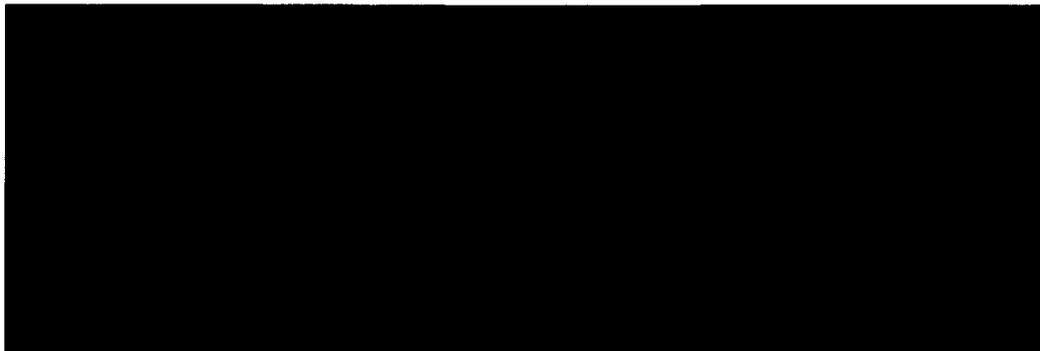


Figure 4.13

Figure 4.14

Figure 4.15

Per-attempt cost calculations yield significantly different results from the DBG-style cost, however. Figures 4.13 through 4.15 show the results of three runs with increasing correlation coefficients. Even with coefficients as high as ± 0.95 the polymorphism continues to have a persistent basin of attraction. Although the basin of attraction does diminish in size as correlation increases, its size remains significant for reasonable degrees of correlation. Also, the equilibrium shift for the polymorphism has returned using per-attempt cost calculations, since costs are incurred both for individuals using the modest strategy as well as those using the greedy strategy.

The primary reason for the persistence of the polymorphic basin with per-attempt cost calculations is that search costs increase exponentially with decreasing proportions. This makes fair individuals who are paired with other fair individuals in the lower regions of the triangle (where fair individuals are scarce), pay a significant search cost penalty on average. Greedy individuals that successfully pair in this same region have far smaller costs in general. Again, per-attempt cost was intended also to model the advantage greedy individuals had in terms of searching duration due to their increased potential payoff. In the region of the polymorphic basin of attraction, this advantage almost always translates into a successful pairing with a modest individual, while fair individuals often fail to pair.

4.3 Synopsis of Chapter IV

This chapter has reviewed many different models, and explored many different variations and enhancements of Skyrms' original divide-the-cake model. It is not entirely clear at this point whether the model actually retains the polymorphism basin of attraction under realistic assumptions or not. Altering the correlation scheme to make it more plausible gives different results depending on how the scheme is actually altered. However, the most plausible scheme from an individual's view (one that maximizes individual advantage, namely BEFS-style) has the basin disappear, but only at very large correlations. Whether such large correlations would actually develop in real situations is still very much an open question, which I will examine in the next chapter.

The addition of cost calculations poses other difficulties as well. While D'Arms, Batterman, and Gorny's original cost calculation method increased the size of the polymorphic basin even further than before under their own correlation scheme, it had certain shortcomings which made models using it incompatible with earlier models, forcing it to need revision. The revised cost calculations performed essentially the same way as the original in the DBG-style model, increasing the polymorphism basin size. However, it failed to do so under the

improved correlation scheme in the BEFS-model, and the polymorphism basin once again disappeared.

An alternate method of search cost calculations, per-attempt cost, attempted to correct some of the shortcomings of the both the original and revised cost methods, such as penalizing strategies even under extremely advantageous conditions, failing to scale the cost appropriately under increasingly difficult situations, and neglecting the advantages of the greedy strategy in terms of cost paying ability. Per-attempt cost gave smaller basins of attraction for the polymorphism than the other cost methods, but those basins persisted even in cases of extremely high correlations. Per-attempt cost is itself limited in some regards, and even more work could be done towards increasing the realism of search cost calculations in these models. Still, the results here do indicate that search cost factors by themselves are enough to change the behaviour of the model significantly, quite separately of the correlation scheme used.

Chapter V - Correlation

5.1 Re-examining Correlation

Up until now we have been exploring models with increasing complexity in the correlation scheme: from the simpler self-correlation of Skyrms' model, to the strategy dependent scheme in the DBG-style model, to the 'enemy-specific' anti-correlation in the BEFS-style model. Each of these changes was described as an increase in the realism of the model by those proposing them. Furthermore, each of these increases in realism corresponded to making correlation more individually advantageous. For example, the move to strategy-dependent correlation for greedy individuals in the DBG-style model was based on the observation that individuals using the greedy strategy have good reason *not* to seek out other individuals using the same strategy. It was further refined in the BEFS-model by going one step farther and recognizing that individuals using the greedy strategy have just as much motivation for avoiding fair individuals as they do for avoiding themselves.

Every model examined thus far has one element in common: correlation is constant for all individuals using a particular strategy. Correlation in some sense has been globally imposed on individuals, although we have been steadily making that correlation more advantageous to the individual it is imposed on. All the arguments that have been looked at so far concerning correlation are based in terms of what *type* of correlation an individual would prefer, i.e. correlation for or against certain other strategies. We have neglected to examine how *much* correlation an individual might prefer. Take Skyrms' original model, for instance. In it, we observe that the polymorphism basin disappears for correlation values of 0.2 and up. But why should we believe that 0.2 is more realistic than, say, 0.5 or 0.01 or any other number?

To investigate this kind of question, we need to begin by looking at the model in a new way. Rather than seeing correlation values as simply an aspect of the model, we instead treat correlation as the problem to be solved. Previous

models used a given correlation scheme and fixed correlation values and found the equilibria and basins of attraction that resulted for every different starting point. Now, we re-examine those models and seek to determine what correlations develop for different starting points rather than what equilibria develop. To do this, we take correlation and tie it to the individual rather than to the model or the strategy. Then we let the dynamics of the system find the correlation values based on how well the individuals do in pairing. Just as previous models renormalized a population based on how successful a given strategy was, we now renormalize based on how effective the correlation values are.

5.2 Evolution of Correlation

Updating the models in the upcoming simulations to allow for this type of flexible correlation requires several modifications. First, correlation is now no longer tied to a given strategy, nor is it a uniform value. Every individual will have a correlation coefficient associated with it, each between -1.0 to 1.0 inclusive. These correlation values represent the correlation or anticorrelation that a specific individual has towards individuals following one of the 3 strategy types -- which one depends on the correlation scheme used in the particular model. For example, in a BEFS-style model, the coefficient for a fair individual is interpreted as an affinity or repulsion towards greedy individuals, whereas in a DBG-style or Skyrms-style model it would be interpreted as affinity or repulsion towards other fair individuals.

Another change that needs to be made is in the method of renormalization. All previous models based subsequent populations only on the total payoffs of the three strategies in the preceding population. However, since individuals now have the correlation coefficients associated with them, individuals need to persist from one generation to another, and need to pass on their coefficients to new individuals in cases where the proportion of a strategy is increasing. This is the kind of simulation the code base discussed in Chapter IV was designed to be able to readily accommodate.

All the simulations in this chapter use a different procedure from those outlined in section 4.2.1 and subsequent sections. The new procedure is as follows:

- Initialize a population with desired mix of modest, fair, and greedy individuals.
- Randomly initialize each individual's correlation constant e with a value between -1.0 to 1.0, except for modest individuals in a DBG or BEFS-style model, in which case e is initialized to 0.0
- Select one of the individuals from the population at random and remove it from the population
- Let S_1 , S_2 , and S_3 be the proportion of each of the strategies present in the resulting population (after removal) such that S_1 is the proportion of individuals with the target strategy of the removed individual (see Figure 5.1 below for details), and S_2 , S_3 are the proportions of individuals using the other strategies. $S_1 + S_2 + S_3 = 1$.
- Let e be the correlation factor of the removed individual. Let $S_1' = S_1 + e(S_2 + S_3)$ if $e \geq 0$, otherwise $S_1' = (1+e)S_1$. Let $S_2' = (1 - S_1')(S_2 / (S_2 + S_3))$, and $S_3' = (1 - S_1')(S_3 / (S_2 + S_3))$.
- Choose a random number r such that $r \geq 0.0$ and $r < 1.0$. If $r \leq S_1'$ remove a second individual from the population with the target strategy of the individual removed, otherwise if $S_1' < r \leq S_1' + S_2'$ select a second individual from the population which uses the strategy corresponding to S_2 , otherwise select a second individual from the population which uses the strategy corresponding to S_3 .
- Pair the two individuals removed together in the divide-the-cake game and assign payoff accordingly.
- Repeat the previous 5 steps so long as there are two or more individuals remaining in the population.
- Let M , F , and G be the total payoffs received by all individuals which

played modest, fair, and greedy, respectively. Let N_M , N_F , and N_G be the total number of modest, fair, and greedy individuals in the original population (before any individuals were removed), and $N = N_M + N_F + N_G$. Now let $N_M' = [M / (M+F+G)] \cdot N$, $N_F' = [F / (M+F+G)] \cdot N$, and $N_G' = [G / (M+F+G)] \cdot N$.

- Randomly select N_M' individuals that use the modest strategy and pass them to the subsequent generation. If there are insufficient modest individuals (i.e. $N_M' > N_M$) then repeatedly ($N_M' - N_M$ times) select a random modest individual and duplicate it. Repeat this process for the fair and greedy individuals.

Target Strategies (M=Modest, F=Fair, G=Greedy)

		Model Type		
		Skyrms	DBG	BEFS
Strategy Used	M	M	N/A	N/A
	F	F	F	G
	G	G	G	M

Figure 5.1

5.2.1 Results with DBG-Style Model

I ran simulations using this new procedure on both DBG and BEFS-style models. I omitted Skyrms-style models largely because the polymorphic basin of attraction is eliminated with even modest amounts of correlation, while the other correlation schemes have basins which persist even in the presence of some correlation or anti-correlation. It is the non-Skyrms models for which investigating the magnitude of correlation will be most informative. The resulting basins for the DBG-style run are shown below in Figure 5.2.

Results for DBG-Style Model with Random Individual Correlation

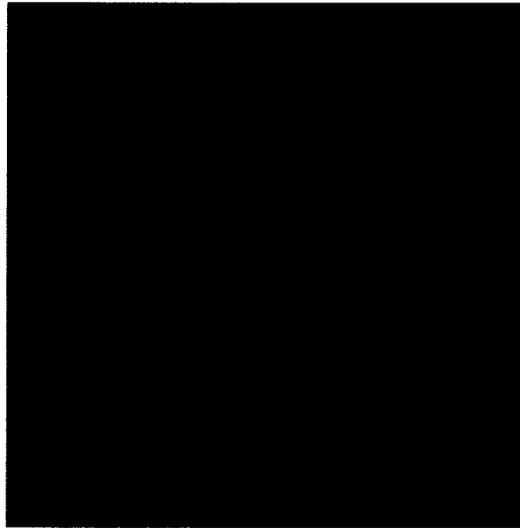


Figure 5.2

The polymorphic basin has almost disappeared in Figure 5.2. This is largely the result of the correlations not being ‘primed’ sufficiently. If individuals following the greedy strategy in the lower regions of the triangle (i.e. populations with low initial proportions of fair individuals) do not initially have a fair degree of anticorrelation, they pair with one another too often. The resulting drop in numbers makes the greedy strategy unlikely in many cases to overtake the fair strategy. Ultimately, the greedy strategy is more dependent on initially advantageous correlation for its survival than the fair strategy is.

The resulting correlations from the DBG-style run are shown below in Figures 5.3 and Figure 5.4. These figures represent a different type of information than previous diagrams, and use a different method of representation. The images show whether the average correlation was positive (blue) or negative (red) for a particular strategy for each starting proportion mix. Since some strategies failed to survive in some areas, those regions of the triangular map are very dark grey, indicating an extinction of the strategy in the final state for every run. For example, in the extreme upper regions of the triangle for Figure 5.3, corresponding to populations which begin with a very high proportion of fair individuals, no population had any remaining greedy individuals in existence after

even a few generations. The simulations were ended after 100 generations, and every population had evolved to equilibrium within this time frame.

In a majority of cases the final average correlation was very small, indicating that correlation played no significant role. This often occurred because convergence to one equilibrium happened too rapidly for much change in correlation values. These are represented in the figures as light grey areas, and the cutoff values are indicated. For example, a cutoff value of 0.01 indicates that all points where the average final correlation was between -0.01 and 0.01 is represented as light grey. These cutoff values were determined somewhat arbitrarily, and were chosen primarily so as to be the lowest value which produced clear and unambiguous results. Also included is Figure 5.2 again, so comparisons can be made between the equilibria basins and the regions for which certain correlations developed.

Average Final Correlation for DBG-Style Model with

Random Initial Individual Correlation

Greedy (Cutoff 0.01)

Fair (Cutoff 0.04)

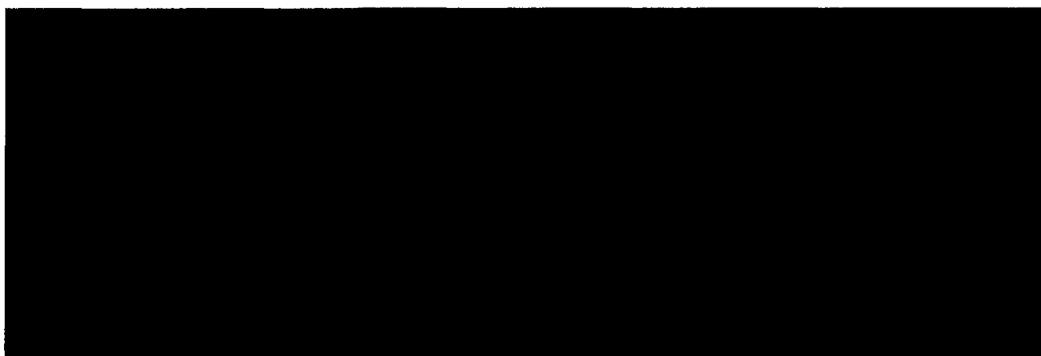


Figure 5.2

Figure 5.3

Figure 5.4

The figures indicate that individuals using the fair strategy in most cases did not develop a significant amount of correlation, as can be seen in the very large light grey expanse in Figure 5.4. In fact, many populations which always ended in the all-fair equilibrium actually had a small but negative average correlation even in regions quite close to the bottom, which is part of the reason a

cutoff of 0.04 was selected for that image. This in no way indicates the actual development of a negative correlation. Instead, roughly half of all populations started out with very small negative average correlation by random chance, and many of those proceeded to the all-fair equilibrium quickly enough that this average had no time to shift to positive values. This same effect is seen for the greedy strategy as well, in the very small blue region at the bottom of Figure 5.3. Here, fair individuals are very scarce and often died out too rapidly, so greedy individuals never needed to shift to negative correlation values.

The real areas where much movement in correlation can be seen are in the extreme lower edge, and also in the lower left corner (with high initial proportions of modest individuals). The lower edge saw many cases where fair individuals with high positive correlations survived for 100 generations, sometimes outcompeting the greedy individuals which greatly outnumbered them initially. Most of the populations in which the greedy/modest polymorphism became dominant did so largely on the basis of overwhelming the fair minority, rather than on the basis of correlation. The exception to this was in the lower left corner, where the greedy strategy did consistently develop negative correlation. Even where a given population failed to end in the modest/greedy equilibrium, the development of negative correlation often allowed the greedy strategy to persist for a significant number of generations.

One more final note concerns the very tiny light grey dot very high up in Figure 5.3. This represents a single population which managed to have the greedy strategy survive right up to the final generation. Actually, only one individual remained out of a population of 750, but this still demonstrates that the introduction of evolutionary randomness and finite populations can from time to time bring forward some rather implausible results due to sheer chance.

The final average correlation for greedy individuals where they actually survived 100 generations was typically between +0.03 and -0.03. The most extreme correlations developed in populations with high initial proportions of modest individuals, and this region held some of the longest struggles between

greedy individuals with high negative correlation and fair individuals with high positive correlation. Negative correlation for greedy individuals sometimes ranged as high as -0.31 at generation 100, and fair individuals in this same region often had positive correlation above 0.5. One good illustration of this kind of prolonged competition was one population (initially 88% modest, 4% fair, 8% greedy) where greedy individuals by generation 19 had achieved -0.25 correlation, increased their proportion to over half the population, and reduced the fair proportion to 2.4%, only to be overtaken when the surviving fair individuals surpassed 0.62 correlation and drove the greedy individuals to extinction 60 generations later (despite greedy ultimately achieving -0.39 average correlation). Occasionally, the last few surviving fair or greedy individuals before extinction might have extreme correlation values close to ± 1.0 , but the diagrams above ignore these extremes and focus instead only the final averages at generation 100.

These simulations do not prove that certain correlation values are necessarily more or less realistic than others, but can be used to provide some kind of starting point for discussions. In general, it seems that in DBG-style models correlation in many places is nearly irrelevant to the outcome. Most starting points in the population space rapidly converge to one or the other of the two equilibria, and final populations can have correlation values that are even somewhat detrimental. Also, final values are almost always under ± 0.6 no matter how long the struggle lasts, and can provide a measure of what constitutes 'realistic' correlation rates. For example, if we think of the populations as being initially formed by mixing some existing populations in the desired ratios, those existing populations might well have had correlation values in this range as a result of previous encounters. However, the fact that some few persistent individuals in very prolonged struggles can achieve more extreme correlation values (for a generation or two) could be used to argue that unlimited correlations should be considered. For instance, if populations are instead formed by migration of individuals crowded out of nearby regions (due to extreme competition) to some new unpopulated region, then very highly correlated

individuals could be represented.

5.2.2 Results with BEFS-Style Model¹²

The results of this same kind of simulation (i.e. using random initial individual correlation) on a BEFS-style model produced very different results from those in the DBG-style model. The resulting basins of attraction are shown in Figure 5.5 below.

Results for BEFS-Style Model with Random Individual Correlation

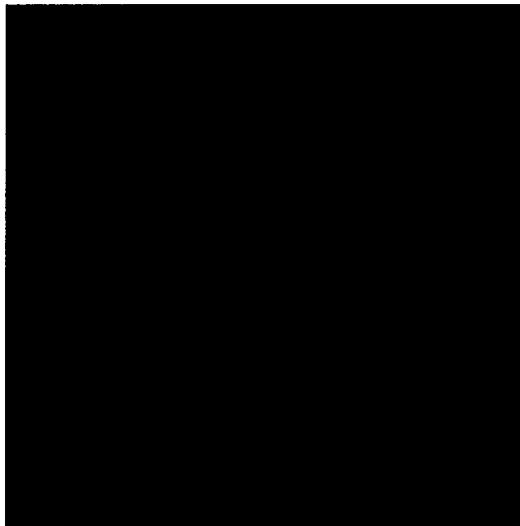


Figure 5.5

There are some very significant differences between this image and the same one for the DBG-style model (Figure 5.2), particularly in the size of the region dominated by the modest/greedy polymorphism. However, a careful observation of the Figure 5.5 will reveal that this area is heavily 'speckled,' which indicates that even populations with initially very low proportions of fair individuals quite often end up in the all-fair equilibrium anyway. This region has highly variable behaviour, and almost every starting proportion mix in the lower region had at least one population which managed to reach the all-fair equilibrium rather than the modest/greedy equilibrium. The correlation results are shown below.

*Average Final Correlation for BEFS-Style Model with
Random Initial Individual Correlation*

Fair (Cutoff 0.02)

Greedy (Cutoff 0.0001)

Greedy (Cutoff 0.01)

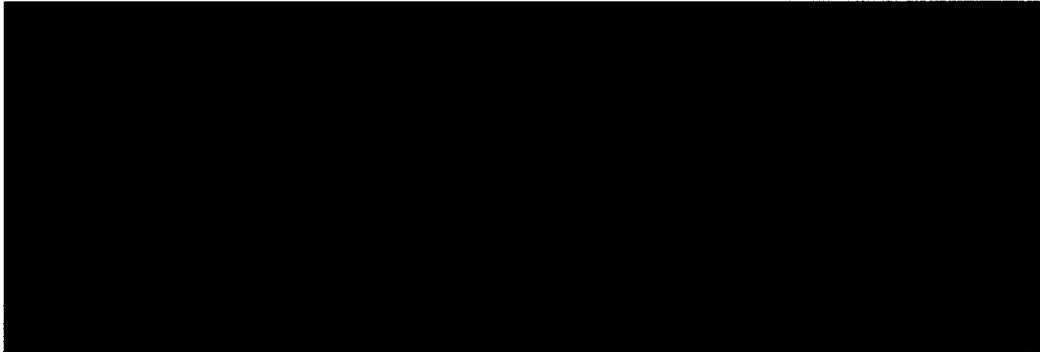


Figure 5.6

Figure 5.7

Figure 5.8

Recall that in the BEFS-style model, fair individuals use their correlation to avoid greedy individuals, so we should expect to see negative correlation (red areas) for fair. Similarly, greedy individuals in the BEFS-style model are correlated to modest individuals, so we should see positive correlation (blue areas) for the greedy strategy. Indeed, fair individuals did develop negative correlation over a large area, and even in the area dominated by the modest/greedy polymorphism there are remarkably few initial starting positions which did not end up in the all-fair equilibrium on at least one run. Individuals following the greedy strategy, however, very often ended up negatively correlated rather than positively as expected, except in a small band at the edge of the basin, clearly seen in Figure 5.7. However, the positive correlations in this blue band were typically very low in magnitude, and Figure 5.8 with a higher cutoff value reveals that most correlations in this region were less than 0.01 except in an area corresponding to high initial proportions of modest individuals. This same region also produced some of the most prolonged competitions (slower convergence) and most extreme correlations for the DBG-style model as well.

In the upper region of the triangle (light grey region in Figure 5.6, or dark grey region in Figure 5.7), fair individuals rapidly overtake individuals using the other strategies. This occurs quickly enough that the average correlation remains

unchanged, and so the average correlation is very close to 0.0 in this region.

In the lower regions there was much more variability in the possible outcomes. Some populations would lead to the all-fair equilibrium, occasionally as the result of *a single* fair individual with extreme anti-correlation. How extreme the anticorrelation needed to be in order for this to happen depended very greatly on the average greedy correlation. For instance, in one population a single fair individual with correlation -0.64 at generation 11 had taken over the entire population by generation 76, while in another population a fair individual with correlation -0.97 failed to survive more than a couple of generations. The cutoff point seems to be roughly -0.40 correlation; lone fair individuals with correlation closer to zero than this never managed to take over the entire population, although some managed to survive a substantial amount of time. Due to this possibility of invasion by even a single fair individual, we have the ‘speckled’ appearance of the lower region in Figure 5.5 earlier.

One important aspect of this simulation is that even after reaching what would have been an equilibrium in the other models, when either all fair individuals are gone or all greedy individuals are gone, there are still pressures to change the correlation over time. This explains some of the more puzzling behaviour seen above, such as why greedy individuals have negative average correlation over such a large portion of the triangle, when the reasoning used in section 3.2 suggested positive correlations would be the most individually advantageous. In cases where fair individuals became dominant, their correlation remain unchanged from that point on, hence the substantial red region in Figure 5.6 as we would expect. However, when the fair individuals were driven to extinction in the population, the equilibrium tended to continue to evolve due to pressures from the overshoot problem discussed in section 4.2.4 earlier. While greedy individuals might obtain a fairly high positive correlation towards modest individuals while competing with fair individuals, the correlation tended to approach zero once again after they were gone. This pressure occurred most strongly when the average correlation was positive, and less so when the average

correlation was negative. The result of this is that the final average correlation was negative except in the region where competition pushed the greedy correlation values sufficiently high that this levelling drift did not have time to work. This is the blue region of Figure 5.7. This tendency of the greedy correlation to level off is also the reason why the overshoot red-shift in basin colour seen in Figure 4.8 through 4.10 is not observed in Figure 5.5: small or negative correlations tend to balance the modest to greedy ratio and make it ≥ 1 .

Note that this correlation drift in the modest/greedy polymorphic basin is more or less an artifact of my simulation method. If the simulation for a given population had been stopped when fair individuals reached extinction, rather than waiting for the correlations to settle down, the entire non-grey region of Figure 5.7 would have been solid blue. Also, one might expect that since greedy individuals with very high correlations are more likely to pair successfully with a modest individual, they would be more likely to be represented in the subsequent generation, tending to drive the average correlation higher rather than down.

However, the method I used to determine the subsequent generation was based on random selection instead of success at pairing. If this happened to produce (though sheer chance) a greedy population with lower average correlation, this would decrease the greedy to modest ratio of the next generation which in turn tended to increase the overall total fitness of the greedy strategy. This increase of the total fitness entails the generation following that would have more greedy members, created by randomly selecting and cloning individuals, which in this case were the individuals with lower average correlation. Effectively, this gave lower correlation greedy individuals a reproductive advantage of sorts. Drift in the opposite direction also occurred, but primarily because highly anti-correlated individuals tend to drop the total payoff of the greedy strategy as well. However, as mentioned earlier, this was a slower process in general.

If I had used a slightly more realistic method of constructing a subsequent generation, this would have resulted in a tendency for greedy correlation to climb

close to 1.0 or perfect correlation. This would result in a higher greedy to modest ratio than in the uncorrelated case. Interestingly, such a population is much less resistant to invasion by fair individuals, particularly if they have a very negative correlation. This might lend support to Skyrms' claim that given enough perturbation of the system, we can expect the vast majority of the time to be spent in the all-fair state.

In the BEFS-style model, there seems to be little reason to limit the choice of correlation coefficient to some particular range. Populations of fair individuals in particular often had extremely high negative correlations of -0.80 and beyond. Although the levelling drift tended to even out the greedy correlations, at the point where fair individuals became extinct the correlation was often above 0.80. Moreover, it can be convincingly argued that a better proportion growth model (i.e. one based on individual fitness) would tend to drive greedy correlation values even higher. My opinion is that when looking at BEFS-style models, extremely high correlation coefficients close to ± 1.0 cannot be ruled implausible.

Chapter VI - Multi-level (Group) Selection

6.1 Group Selection and Altruism

If every model we had considered in Chapter IV had only the all-fair equilibrium retaining a substantial basin of attraction, we might consider the question concerning the viability of Skyrms' model closed. Increasing the realism of the model by changing the correlation scheme as in the DBG and BEFS-style models would not discount the model if subsequent improvements still yielded a global all-fair equilibrium, and we would still have a generalist evolutionary argument that could bypass the Universality and Possibility problems.

Unfortunately this did not quite happen. The investigation in Chapter V revealed that almost any correlation values could conceivably evolve in a population, and even extremely high correlation values cannot be easily excluded as impossible. However, even the most extreme correlation values do not suffice to remove the basin of attraction for the modest/greedy polymorphism in some models, particularly in variations of the BEFS-style model (with or without cost calculations). Moving the correlation coefficients upwards does not by itself resolve the issue of the viability of the model.

There are other evolutionary paradigms that we can look into though, particularly group selection models. Essentially, group selection looks at evolution as a composite of two separate kinds of selection: within-group and between-group. A strategy or trait that might have lower fitness under within-group selection can nonetheless thrive under some circumstances if that strategy or trait improves the fitness of the group as a whole for between-group selection.

The notion of group selection historically has received some bad press, and spent a considerable amount of time as a largely rejected theory. Despite being initially proposed by Darwin¹³, group selection is still considered by some authors as being an unlikely or unnecessary theory even today. The book *Unto Others* by Elliott Sober and David Sloan Wilson (Sober, & Wilson, 1998) gives an excellent historical discussion of the history of the controversy, and provides argument that

the group selection paradigm is viable and in fact does occur in nature.

There are other alternatives to the group selection view. It is entirely possible to look at biological evolution, for instance, as acting strictly at the level of the gene and avoid any reference to groups at all (Sober, & Wilson, 1994, p. 536). Brian Skyrms in particular looks at group selection as being simply a different way to talk about correlation (Skyrms, 2002, p. 699). Sober and Wilson agree that viewing group selection in these terms is valid, but argue that focusing on correlation alone can lead to erroneous conclusions, particularly when looking at evolution in situations where interaction between neighbours is governed by a fixed spatial structure (Sober, & Wilson, 2002, pp. 720-1). In most cases, whether to view evolution in terms of correlation, genes, or group selection is largely a matter of semantics, and which is preferable depends largely on what kind of activity we want to draw attention to (Maynard Smith, 1998, p. 639-40).

For the models that we have been using so far, thinking in terms of group selection or in terms of correlation is more or less equivalent. I have used the term 'correlation scheme' throughout this thesis up to this point, differentiating the different models in terms of how the correlation functioned. However, I could just as easily have discussed them as group selection models. For instance, in the models of Chapter V the three strategies would each correspond to a different group, with between-group selection being implemented based on the total fitness (i.e. the renormalization step), and within-group selection being based on an individual's correlation coefficient (i.e. how likely a given individual is to continue on to subsequent generations). Even though the models of Chapter IV did not have such an obvious within-group component, they could still be thought of as group selection models depending how such models are defined (Sober, & Wilson, 1994, p. 536).

Thinking of the models in terms of group selection or in terms of correlation are in many ways equivalent, but I am going to begin discussing the models that appear in the next chapter largely in terms of group selection. Mostly this is to increase clarity, since the upcoming models are going to look at

evolutionary scenarios in which multiple populations coexist, each containing a mix of the three strategies. Rather than discussing hierarchies of correlation schemes and doing complex calculations to find the average correlations of a strategy (which will typically vary from population to population anyway), I will be referring to each population as a group. This makes explicit the tension between the fair individuals and the modest/greedy individuals within a given population (or group), and also allows the competitive performance of one population versus the others to be discussed separately.

Group selection is typically used to explain the evolution of altruism. Altruistic behaviour can be thought of as behaviour that is disadvantageous to the individual, but advantageous to the group. A common example is giving a warning cry when spotting a predator: giving the call draws attention to the individual making it (making it more likely to be attacked), but increases the survival chances of the group as a whole by allowing them to avoid the predator. Groups that have a high proportion of individuals giving warning cries will do better overall than other groups, and this allows the continuation of the altruistic trait despite its being maladaptive for the individual.

There are several similarities between altruism and the fair strategy in Skyrms' divide-the-cake model. For example, groups containing a large proportion of fair individuals will usually outperform groups with none (i.e. all-fair groups have a higher total fitness). Also, being fair can sometimes be disadvantageous to the individual, particularly when greatly outnumbered. The biggest difference, of course, is that this disadvantage depends greatly on the current proportions of strategies within the group, and the fair strategy can just as easily become advantageous as the proportions change. Whether fairness is a type of altruistic behaviour is at best very questionable and, like many social traits, would depend on the level of selection one looks at (Hamilton, 1975, p. 135). Since fairness is not disadvantageous to the individual in the majority of cases, it probably does not qualify as altruistic under most definitions, and I will not attempt to classify it as such. Nonetheless, the same kinds of group selection

models that are commonly used when investigating the evolution of altruism might still be worth looking at in the context of fair versus modest/greedy in Skyrms' models.

6.2 Wright's Trickle Model

Wright's trickle model was developed by Sewall Wright in 1945, and is one of the simplest group selection models. In it, different populations are isolated from one another, except for a very small amount of dispersers or migrators who 'trickle' from one population to another in each generation. This model is also sometimes called Wright's island model. Figure 6.1 shows a graphical representation of this model with five populations or islands.

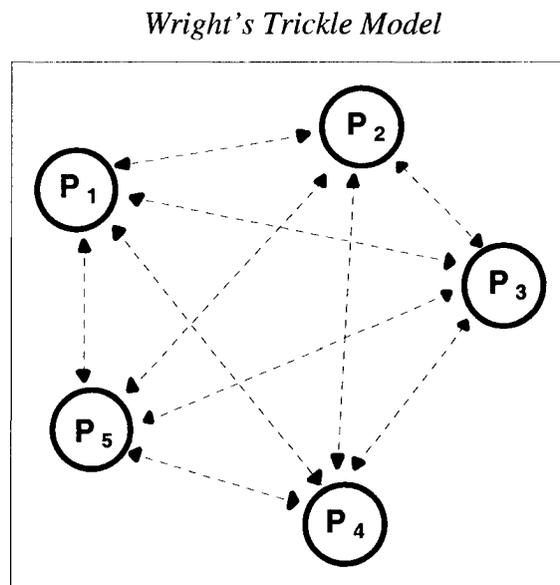


Figure 6.1

The acting principle in Wright's trickle model is genetic drift. Genetic drift is a change in the frequency of a strategy or trait due to random sampling effects between successive generations. That is, traits can vary randomly in frequency due to statistical variation alone. For example, genetic drift will occur when two strategies or traits in a population are equally fit and well represented,

and cause their proportions to vary randomly until eventually one or the other will become dominant while the other becomes extinct. Even though altruism is more likely to diminish in frequency due to natural selection, genetic drift can nevertheless cause it to become dominant despite its disadvantage, particularly in smaller populations. Ideally in Wright's model, altruism should become dominant in a few populations via genetic drift, and then these altruistic populations will out-compete the other groups by lasting longer and producing more dispersers.

Wright's model ends up doing a fairly poor job in general of evolving altruism (Sober, & Wilson, 1998, p. 61). Altruistic populations are often poorly resistant to invasion initially, and incoming non-altruistic individuals cause most populations to drift away from altruism. One solution is to reduce the trickle of migrators, but this in turn reduces the ability of altruistic populations to out-compete others. Another solution involves having most of the migrators found new populations rather than simply enter existing populations. This can be made to work, but only if being altruistic provides an advantage either by increasing the number of migrators, or by making new colonies more resistant to extinction (Hamilton, 1975, pp. 141-2).

Consider the situation in which the populations or islands in Wright's model are replaced with populations from Skyrms' model. All-fair populations in Skyrms' model (or DBG or BEFS-style models) are significantly better adapted to survival within Wright's trickle model than the altruistic populations previously described are. Most populations in Skyrms' model are either highly resistant to invasion initially, or become so within a small number of generations. Also, Skyrms' model already has proportion growth based on total fitness, and this is readily adapted to increasing the numbers of migrators. Although the altruistic populations in Wright's trickle model become fixed through genetic drift, this is not required by the populations of Skyrms' model since they are quickly carried to one equilibrium or the other by the characteristics of the model (i.e. renormalization). That is, populations in Skyrms' model frequently fall into the all-fair equilibrium, and do so via natural selection alone. Since genetic drift is

not a requirement for fixation when applying the trickle scenario to Skyrms' model, we can retain the generalist nature of Skyrms' model and still apply it equally well to either biological evolution or social/cultural evolution. These characteristics make the trickle model a promising candidate for extending Skyrms' model.

6.3 Haystack Model

Another group selection model often used in discussions of altruism is John Maynard Smith's haystack model, so named because it describes a whimsical species of mouse that lives its entire life in haystacks (Maynard Smith, 1964). This model progresses in two discrete periods: a reproductive phase, and a dispersal phase. The reproductive phase proceeds for some number of generations, during which individuals reproduce entirely within their own population. The reproductive phase is followed by the dispersal phase, in which all the individuals of every haystack are pooled into a global population, and then disbanded randomly into new groups. This cycle of alternating reproductive and dispersal phases is repeated several times.

Like the trickle model, some populations would become altruistic through genetic drift (but this can be dropped in Skyrms' case as it was for the trickle model). These altruistic populations would out-compete the others by producing more dispersers as a result of having a larger population size. A graphical representation of this model is shown below in Figure 6.2, with altruistic populations having a different pattern from non-altruistic populations.

Maynard Smith's Haystack Model

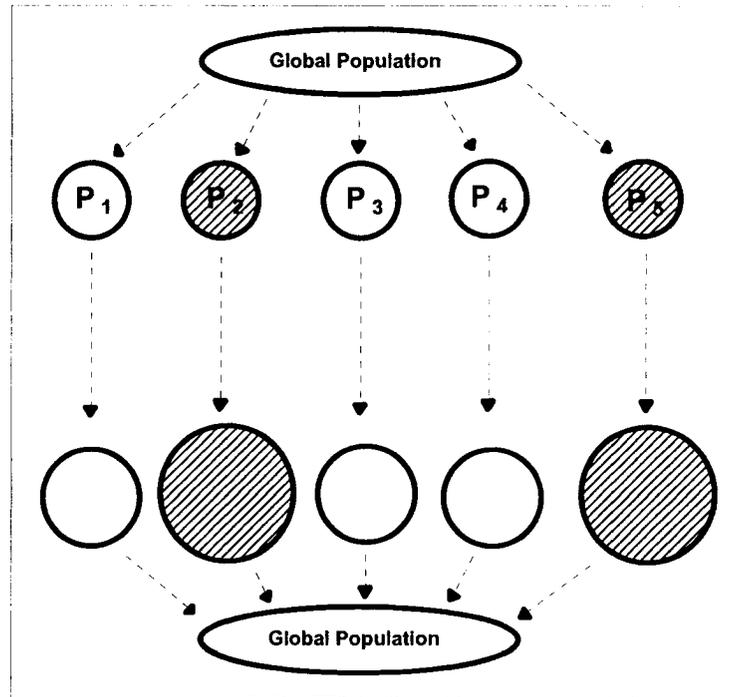


Figure 6.2

Maynard Smith's original haystack model has each population founded by a single individual. This would not be compatible with Skyrms' model which relies on some initial mix of strategy types and cannot be generalized to single individuals. In *On the Founder Effect and the Evolution of Altruistic Traits* (Cohen, & Eschel, 1976), several extensions of the haystack model are considered which are much closer to the kind of framework needed for Skyrms' model. In particular, I will be extrapolating from the 'case I' framework in which groups are founded by some fixed number of individuals, and every individual has an equal chance of being included in each founding group independently of the types of other founding individuals.¹⁴

The haystack model has better success than the trickle model in allowing the evolution of altruism. Although altruistic individuals typically decline in terms of proportion within any given population during the reproductive phase, the advantage they confer to the group allows their numbers to grow in absolute

terms. This in turn allows altruistic individuals to increase in proportion when pooled together during the dispersal phase.

The fair individuals within populations in Skyrms' model do better than altruistic individuals in the basic haystack model. In some cases, fair individuals do decline in proportion during successive generations (eventually becoming extinct), but in many cases they instead increase in proportion and quickly become dominant driving the other strategies to extinction. If populations in the all-fair equilibrium can turn their higher efficiency into a reproductive advantage, the increased proportion of fairs ought eventually to create initial populations (i.e. after a dispersal phase) well-situated to developing the all-fair equilibrium. These in turn will then outperform other neighbouring populations which develop the greedy/modest equilibrium instead.

Chapter VII - Coding Project Part II

7.1 Description of the Group Selection Coding Project

In order to examine how the group selection models of the preceding chapter will affect the behaviour of the models, I have extended the code to incorporate multiple populations evolving concurrently. Most of the discussion for the various correlation schemes in Chapter IV remains unchanged in these newer models. Any given population running a specific correlation scheme (such as DBG-style) still uses the same steps in the same order as outlined previously, and any additional interactions between separate populations occur in addition to those original steps rather than replacing them. Upcoming sections will describe any additional steps or processing used.

Since we now have multiple populations, and those populations are interacting with one another, this made the simulations run much slower than before. As a result, I had to cut down the number of separate runs which were averaged together for any given starting point. This did not significantly reduce the accuracy, since the number of runs averaged together in Chapter IV was considerably higher than what was actually required.

One other consideration in my implementation concerns how the colour diagrams are produced. Since there are multiple populations running concurrently now, I needed to determine which of those populations will be represented. I solved this by selecting a single population from each run, treating it as the 'target' population. The result images were produced by averaging the resulting final state for the target population over several different runs. This was simpler than other methods, such as presenting separate diagrams for every population. Averaging all the populations together rather than selecting a target population would also not have worked -- particularly for the trickle model for reasons explained in the next section. Otherwise, the diagrams are produced in the same manner and using the same conventions as before.

The group selection models also abandon the fixed population size

assumption of the original models. The trickle migration step and haystack dispersal phase require populations to temporarily change size while individuals move to or from a population, so any given population is allowed to either grow or shrink as necessary. As stated previously, this is accomplished as an additional step to the original models rather than replacing the normal renormalization step, and any population size increase or decrease is done only after renormalization. Since population size decreases are deliberate in this model, any populations which have total fitness of zero do *not* get renormalized to a single greedy individual in these group selection scenarios. Should a population have total fitness of zero or somehow lose all remaining individuals, it proceeds as an empty population for the subsequent iterations. Such a population can still accept incoming migrators in the trickle model (and proceed normally from there) or function as a normal destination haystack during the dispersal phase, but otherwise does nothing.

7.2 Trickle Model

My implementation of the trickle model essentially consists of multiple populations each using one of the correlation schemes described earlier (i.e. DBG, BEFS, etc.). A generation proceeds as normal, with each separate population doing the usual steps in the same order as before. After the final renormalization step, however, a new migration step is added in which some individuals in each population are moved to a separate migration pool. Each population then sequentially takes a random individual from the migration pool until none remain. As an analogy, consider individuals as playing cards in separate hands. A subset of each hand is randomly selected and removed to a new pile, which is shuffled and dealt back into the original hands. This technically allows an individual to end up back in the population it started from, so in my implementation Figure 6.1 would have additional arrows looping from each population back to itself.

Formally, the model takes the normal steps and applies them to each population individually. Then, the following step is appended to those normally

performed in each iteration:

- For every population $P_1 .. P_N$, randomly select some subset of individuals and remove them to population M . Each population $P_1 .. P_N$ in turn takes one randomly selected individual from M . Repeat until no individuals in M remain.

One difficulty in implementing the trickle model was choosing how the populations should be initialized. Previously, any given point in the triangle corresponded to a single population with a specific initial proportion for the different strategies. However, in the trickle model initializing all the populations to the exactly same strategy proportions would have had no real effect. Since all populations are the same, randomly moving a small number from one identical population to another would have effectively resulted in no change at all. In the vast majority of cases, we would be left with a set of populations performing more or less the same activity as if they were completely isolated.

One of the underlying ideas in Wright's model was that all of the populations should have the capacity to become either altruistic or nonaltruistic. Altruism might be antecedently unlikely, in which case only a minority of the populations would fixate to altruism. Alternately, if altruism were very highly likely, a majority of the populations would fixate to altruism. In either case, however, the populations would nonetheless have representatives of both possibilities. When we transpose this idea to the all-fair vs. modest/greedy situation, it becomes necessary to initialize the populations separately from one another. Otherwise, almost every trial will end up having every population rapidly proceeding to exactly the same equilibrium and becoming identical. To avoid this, I initialize only the target population to have the specified strategy mix. The other populations are randomly and independently initialized, with every possible strategy mix being equally likely.

The results for the completely uncorrelated version of Skyrms' model are shown below in figures 7.1 through 7.3 for different migration rates. The number

of migrating individuals from each population that are pooled and redistributed are shown as a proportion of the total number of individuals. Since the number of individuals in a population remains constant from the end of one iteration to another in my implementation, this could just as easily have been described as 5, 10, and 15 migrators rather than 0.01, 0.02, and 0.03 migration rate, respectively.

Trickle Model using Skyrms' Populations (Uncorrelated)

Rate 0.01

Rate 0.02

Rate 0.03



Figure 7.1

Figure 7.2

Figure 7.3

The results of increasing the trickle rate (migration rate) are quite dramatic. The 1% rate has a noticeable impact already, by allowing some populations to evolve to the all-fair equilibrium despite being well within the normal boundaries of the basin for the modest/greedy polymorphism. Increasing the rate to 2% allows populations in all areas of the triangle to potentially become all-fair, and the populations which evolve to the polymorphism have become a minority, even in those regions which are quite near the bottom of the triangle. By the time the rate hits 3%, the modest/greedy polymorphism is almost completely gone. Although not shown in the above images, a migration rate of 5% resulted in only about 32 modest/greedy populations out of a total of 13,260 runs, and the polymorphism was entirely eliminated from every run by 6%.

These results demonstrate that even a reasonably small amount of migration among populations can have a significant impact. There is a feedback mechanism involved in this process, becoming particularly powerful once a

certain threshold has been passed. Increasing the migration rate increases the number of incoming fair individuals, which in turn increases the number of populations moving towards the all-fair equilibrium, which increases the proportion of fair individuals migrating out.

The number of populations (or islands) also has an impact, particularly for determining the migration rate at which the modest/greedy polymorphism disappears. The above images were created using 12 populations in the migration structure. Increasing the number of populations beyond this had only a minor impact on the results. For example, with 50 populations and a 5% rate only 3 populations out of 13,260 reached the polymorphic equilibrium (rather than the 32 seen with 12 populations). Reducing the number of populations had the reverse effect, and one set of simulations with only 5 populations trading migrators needed a 10% migration rate to completely eliminate the polymorphism.¹⁵ In effect, reducing the number of populations increases the probability that a majority of them might be predominately modest/greedy, thus increasing the number of modest/greedy individuals migrating and encouraging the polymorphic result rather than the all-fair.

How convincing these results are will depend greatly on one's intuitions regarding how plausible the trickle model is, what constitutes a reasonable number of populations sharing migrators, and what a reasonable migration rate is. For example, 10% would not seem to qualify as a 'trickle' in most cases, although 3% might be acceptable. It is difficult to make judgments regarding these assumptions without discussing some particular evolutionary type (such as biological or cultural evolution). However, moving into specific domains like this would sacrifice much of the generality that makes Skyrms' argument useful.

The images above at the very least demonstrate that the correlation inherent in the trickle model *can* potentially be used as a substitute for the other correlation schemes looked at previously in Skyrms' model. One difficulty we looked at earlier was how precisely the 0.2 correlation came about. The discussion in Chapter V showed how it might come about if correlation were an

individual-specific characteristic, but the trickle model provides the first step in exploring an alternate method: correlation introduced by group structure. Again, how convincing this alternate method of producing correlation is depends on one's intuitions concerning the plausibility of group or multi-level selection in actual populations, but there is evidence suggesting that it is a factor in many situations (Sober, & Wilson, 1998, p. 54).

Another interesting result of the trickle model investigation is that it provides a means of having complete 100% domination of the triangle by the all-fair equilibrium. Every other model previously explored had populations that initially contained no fair individuals, and these populations never had any possibility of evolving to the all-fair equilibrium. Skyrms mentions the possibility that a population that falls into the polymorphic trap might eventually bounce out of it if there is sufficient random variation in the evolutionary process (Skyrms, 1996, p. 21). The trickle model explored in this section provides a means of providing precisely the kind of 'bump' required to do this, but without requiring the need to introduce individuals possessing a strategy that does not depend on the previous generation. This might be more acceptable than other methods which would require many individuals to simultaneously 'mutate' to a different strategy, or to have many new individuals suddenly appearing from nowhere.

7.3 Haystack Model

As previously described, my implementation of the haystack model is close to the Cohen-Eschel Case I variant. As in the trickle model, it consists of several populations running concurrently. Rather than a trickle of migrants, however, the haystack model completely dismantles all populations and redistributes the individuals into new ones periodically (called the dispersal phase). During the intervening generations, the populations are completely isolated from one another.

Just as with the trickle model, there exists a difficulty deciding how to initialize the populations appropriately. The solution I used in the case of the

trickle model was initializing a target population with the desired strategy ratio, and initializing the others to random ratios. While this method would still work in the haystack model, the results would have been rather uninteresting. Since the majority of populations (haystacks) in a given suitably large set would move toward the all-fair equilibrium during the reproductive phase, the subsequent dispersal phase would result in all haystacks having a majority of fair individuals. The end result would have been every haystack becoming all-fair shortly after the first post-reproductive dispersal phase. Although it could be argued that this would be a viable method of initialization, I felt the results would have been too trivial to really add much to the discussion. One of the ideas inherent in the haystack model was that despite being at a disadvantage, altruistic behaviour might still do well globally provided that it could give some advantage to the group. Moving this to the all-fair versus greedy/modest framework, I decided a reasonable compromise was to initialize every haystack to be precisely the same. For the regions of the triangle in which we are most interested, this puts fair individuals at a distinct disadvantage.

Since all the populations are initialized precisely the same way, we might expect that all of them would evolve in pretty much the same way. If so, the dispersal phase would result in a set of haystacks more or less identical to the original (pre-dispersal) haystacks, and we would notice little difference in the final result images. However, the move to the haystack model requires that fair individuals provide some sort of advantage to the haystack they reside in. I have incorporated this by allowing the population to grow during the reproductive phase. The amount of population growth per iteration is determined by a linear function of the average fitness of all individuals in the haystack. Populations with a higher proportion of fair individuals have a higher average fitness in general than those which are predominantly modest/greedy, so they will grow faster and provide a greater number of individuals during the dispersal phase.

In formal terms, my implementation of the haystack model uses the normal steps for each population during the reproductive phase. For instance, if each

population is a Skyrms-style population, the steps listed in section 4.2.1 would be performed on each population separately during the reproductive phase. However, after the standard renormalization step, the following step is performed for each population on every iteration to provide population growth. As before, M , F , and G are the total payoffs received by all individuals which played modest, fair, and greedy, respectively, and N is the total size of the population.

- Given a fixed population growth factor $\delta \geq 0$, randomly select one individual from the population and duplicate it (i.e. clone it). Repeat this selection and duplicating process until the population size is $N + \delta \cdot (M+F+G)/N$.

The number δ will limit how fast a population can grow. For instance, since the maximum possible average fitness is 0.5, a value of $\delta = 2.0$ would allow the population (at most) to double each iteration during the reproductive phase, while a value of 0.0 keeps the population from growing. Since population is permitted to grow after each iteration during the reproductive phase, population growth is exponential for δ higher than 0. Although populations containing significant proportions of modest or greedy individuals do grow in size, they almost always have an average fitness less than a population with a large proportion of fair individuals, so they do not grow as fast. The difference is generally not too pronounced however until a significant number of iterations have passed. For example, with $\delta = 0.5$ a population in the all-fair equilibrium grows about 7% faster per iteration than one in the modest/greedy polymorphism. Eventually, this exponential population growth could become problematic, but population sizes are ultimately limited by the dispersal phase which initializes haystacks to a fixed size.

The dispersal phase is handled as a separate step which is performed only periodically. Similar to the way migration was implemented in the trickle model, new haystacks are created by putting everyone in a global temporary population, and then randomly selecting and distributing some of them to haystacks in a

manner analogous to dealing out cards. Formally, the step is as follows, with N_0 being a constant representing the initial size of each population after dispersal.

- After the correct number of reproductive iterations, perform the following: For every population $P_1 .. P_N$, remove all individuals to a temporary population M . Then each population $P_1 .. P_N$ in turn takes one randomly selected individual from M . Repeat until each population contains exactly N_0 individuals. Delete any remaining individuals in M .

Haystack Model using Skyrms' Populations (Uncorrelated)

Dispersal 10, $\delta = 0.5$

Dispersal 15, $\delta = 0.5$

Dispersal 25, $\delta = 0.4$

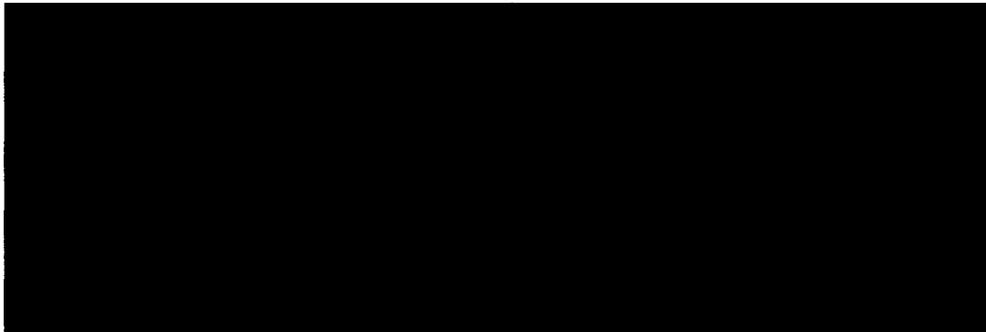


Figure 7.4

Figure 7.5

Figure 7.6

The result images above show the results I obtained for differing numbers of iterations in the reproductive phase. Since all populations (haystacks) were identically initialized and were running identical steps, it was reasonable to average all the final populations together to produce the images above, rather than selecting a single target population as was done for the trickle model. This helped reduce the number of separate runs which needed to be averaged together, which was helpful since the simulations were very slow. Each image is the average of 10 separate runs, each run using 10 haystacks. Also, the population growth factor had to be reduced to increase efficiency for the final image. All simulations were run for only 4 complete reproductive phases, so the number of total generations is increasing along with the length of time between dispersal phases, which is listed

for each image.

In general, there is a noticeable tendency for the all-fair equilibrium to do better as the length of time between dispersal phases increases. This seems somewhat counter-intuitive since the regions of the triangle in which the all-fair equilibrium is encroaching are areas in which the modest/greedy polymorphism is dominant during the reproductive phase, thus reducing the proportion of fair individuals. The longer the populations are isolated, the smaller the proportion of fair individuals that we should expect to survive. In fact, if the reproductive phase is too long, the areas in Figure 7.6 which have switched from the polymorphism to the all-fair equilibrium are areas in which the fair individuals should usually have become extinct after 25 generations, which would in turn contribute to the extinction of fair individuals in other populations after dispersal.

Global extinction of fair individuals was indeed typical for areas in which the fair strategy was a significant minority (proportion of less than 0.11 or so). Below this threshold, most populations quickly become entirely composed of modest and greedy individuals. It seems plausible that if one population should happen to proceed to the all fair equilibrium (and random pairing would allow this to happen once in a while by sheer chance), then the better reproduction rate would allow it to quickly grow larger than the modest/greedy populations and produce populations with more fair individuals after the next dispersal. This was not the most significant factor in basin reduction for the above images, however, since the length of reproductive phase used was too short for growth to make a very large difference. For example, a typical population initialized to 12% fair and 44% modest/greedy that proceeded to the all-fair equilibrium ended up only about 14% larger than the average modest/greedy population did after 25 generations. This is due to the comparatively slow increase in the proportion of fair individuals when greatly outnumbered, resulting in a similar growth rate to other populations for much of the reproductive phase. Even the 25 generations used in Figure 7.6 was not sufficient for many populations moving towards the all-fair equilibrium to finish locally displacing the modest and greedy individuals,

let alone outcompete them globally. However, the reproductive difference did make up some of the added gain for the all-fair basin between Figures 7.5 and 7.6 above.

It turns out the most significant factor in expanding the all-fair basin in the above images is the dispersal phase combined with the fact that the odds of the all-fair equilibrium occurring are higher than the initial fair proportion for that area of the triangle. For example, consider a group of haystacks initialized to 12% fair (I ignore the effects of population growth for now). So long as the odds of such a population becoming all-fair are higher than 12% we should see an increase in the proportion of fair individuals after the first dispersal phase, since over 12% of haystacks would be completely fair. Actually, a rough count shows the likelihood of a population so initialized becoming all-fair is closer to 15% or so. The longer the reproductive phase is, the more likely that populations will reach the all-fair equilibrium and increase in proportion globally after the dispersal phase. The choice of how many haystacks are used will also change how important this factor is. A very large number of haystacks will show an increase in the proportion of fairs in these circumstances most of the time, while a smaller number might not. I used 10 haystacks when generating Figures 7.4 through 7.6, which turned out to be close to the worst-case scenario for the populations shown above, and allowing the fair individuals to increase in proportion only about 46% of the time.

Although not very significant in the images above, population growth could still potentially make a significant impact. The longer the population is isolated, the more pronounced the size differences would become due to the exponential growth difference once the populations had settled into equilibrium. The longer the reproductive phase, the greater the contribution of any population that did manage to reach the all-fair equilibrium. This would be even more pronounced if there were a large number of haystacks contributing to the dispersal phase. If there were over 1000 haystacks, the odds of at least one reaching the all-fair equilibrium are very good for the vast majority of the triangle. If the

reproductive phases were long enough (or the growth factor high enough), that single all-fair haystack would eventually overwhelm all others upon the next dispersal. Simulating such a scenario was simply beyond the capacities of the computers available to me, however, and extremely prolonged exponential growth is unlikely to occur in a majority of real-world populations.

Unlike the trickle model, the haystack model does not seem well-suited to entirely eliminating the modest/greedy polymorphism. However, it still manages to reduce the size of the polymorphism in populations without requiring any sort of correlation other than that produced by the design of the model itself. The model demonstrates that the effects of random pairing in a finite population can be exploited to increase the domain of the all-fair equilibrium without requiring any alteration in the method used to pair one individual with another. Also, the trickle model required other populations to be different from the target proportion mix in order to have a significant impact on the way a population evolved. The haystack model does not have this liability, and can increase the extent of the all-fair basin of attraction even when all populations are initially identical.

7.3 Synopsis of Chapter VII

The results of this chapter are somewhat ambiguous. While the trickle model did succeed in restoring the global all-fair equilibrium, it did so only by adding new populations with different initial starting proportions. This result can be argued to be impractical as a method of establishing the all-fair equilibrium because the global population proportions (which includes all islands simultaneously) is effectively already within the all-fair basin of attraction. That is, if we look at all populations collectively, the initial mix of individuals would naturally proceed to the all-fair equilibrium anyway if only it operated as a single population. A better solution might be to take the required starting proportions and split them into randomly selected subpopulations, so that all islands (taken collectively) would have the desired starting proportion.

Even with the limitations of my trickle model implementation, it

nevertheless manages to demonstrate that even a modest influx of fair individuals can be sufficient to overwhelm the modest/greedy polymorphism equilibrium. Since a 3% migration rate includes 3% migration out of the population as well as in, the net change in fair proportion averages about 0.05 per iteration after equilibriums have been established. This is considerably less than the 0.22 fair proportion normally required initially to guarantee a high likelihood of the all-fair equilibrium in an uncorrelated Skyrms' model. Such a sequence of multiple smaller 'bumps' (randomizations) of the system might be more likely than one very massive one. Also, as mentioned earlier, the trickle model provides these randomizations of the evolutionary process without requiring the need to introduce mutation or entirely new individuals.

There is still a lot of possible exploration that could be done with trickle-style migration models, and what I present here barely scratches the surface. For example, we might argue that random selection is a poor method of selecting migrators, and look at models where migrators are selected based on fitness, with more fit individuals moving more readily. Alternately, models in which strategies with lower fitness migrate more often might be interesting (e.g. strategies are 'crowded out' as suitable niches become scarce). We might also investigate migration in a way analogous to the way Chapter V investigated correlation, making predisposition towards migration a varying individual-specific characteristic and then seeing how it evolves over time.

The haystack model did not entirely remove the modest/greedy polymorphism basin either. The model did manage to reduce it in extent, however, and did so without requiring use of differing haystack proportions -- every haystack was initialized identically. If I had initialized the haystacks to different proportions (by using an initial dispersal phase, for instance), the result would probably have been an even smaller polymorphism basin. Most significantly, the haystack model demonstrated how correlation-like effects might be created using structure alone, rather than requiring correlation to be specified as a fixed population parameter without justification.

The main method by which my haystack implementation increased the size of the all-fair polymorphism was by allowing fair individuals to convert evolutionary randomness into proportion growth. A fairly substantial area of the initial state-space triangle has some significant probability of either equilibrium appearing for a given population. In the case of fair individuals this probability is often higher than the proportion of fairs involved, and the haystack model provides a means of exploiting this. The increase in the all-fair basin would have occurred even if the population growth factor was entirely removed from the model. Figures 7.7 and 7.8 below provide some idea of the maximum change in basin size this effect can produce.

Region for which
Prob(all-fair) > Proportion(initial fair)
(Coloured Blue)

Uncorrelated Skyrms' Model

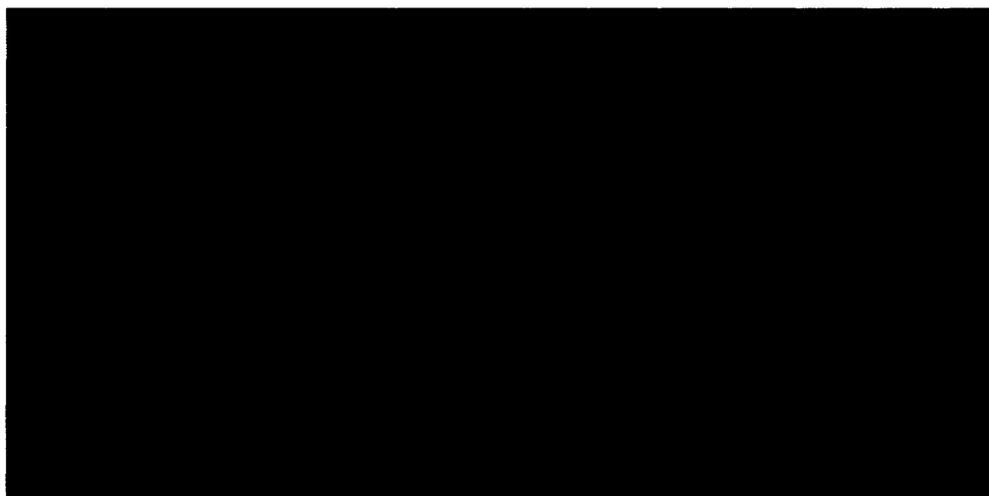


Figure 7.7

Figure 7.8

Figure 7.8 represents the same set of data as Figure 7.7, except instead of being coloured according to the final average proportion of individuals, it has been assigned colours according to whether the proportion of populations ending in the all-fair equilibrium is greater than the initial proportion of fair individuals. The colours for Figure 7.8 were chosen to be similar to the previous diagram in

order to make comparisons more straightforward. It is interesting to note that even though the reduction in size is not very large, the amount by which the proportion of final all-fair populations exceeds the starting fair proportion is at most 0.107 or so, and typically much less. Any changes within a population which make the all-fair equilibrium globally even slightly more likely to occur can thus have a significant impact when combined with a haystack model. For example, Figure 7.9 was produced in the same way as Figure 7.8, except instead of using the uncorrelated Skyrms' model, it shows the approximate region in the case where Skyrms-style correlation is 0.04. This is a very small correlation rate, but when combined with the haystack model it could easily result in the all-fair equilibrium becoming globally dominant again.

*Region in which $Prob(all-fair) > Proportion(initial\ fair)$
for Skyrms' Model (Correlation 0.04)*

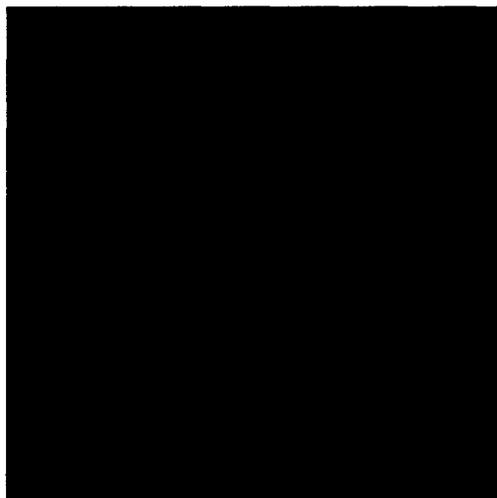


Figure 7.9

As with the trickle model, there remains a lot of potential variations to explore in applying haystack models to Skyrms' model. Variations closer to Cohen-Eschel Case II might be looked at, in which haystacks are initialized by having individuals migrate to random haystacks independently. This would result in non-uniform starting haystack proportions and would likely increase the

domain of the all-fair equilibrium basin even further. Population size increases were not a large factor in the simulations I performed, but variations using higher population growth rates should also increase the all-fair basin under many circumstances.

Chapter VIII - Conclusions and Findings

8.1 Synopsis and Conclusions

I began this thesis hoping to discover whether Skyrms' model could meet the challenges offered by several authors, and attempted to find out whether or not the all-fair basin of attraction could regain its global domination (or a close approximation). I proceeded to look at the adjustments to the model proposed by D'Arms, Batterman, and Gorny, which were intended to make the correlation scheme of the model more realistic. I fine-tuned their model somewhat in turn to increase the realism, culminating in a switch to per-attempt search costs. Although this switch did not eliminate the polymorphic basin, it did reduce its extent to a size comparable to that seen in the uncorrelated Skyrms' model.

I also examined the further extensions to the model proposed by Barrett, Eells, Fitelson, and Sober, which adjusted the correlation scheme further in an attempt to increase realism. Although the BEFS-style correlation scheme actually allowed the all-fair basin to include almost the entire state-space, it did so only for very large correlation values (around ± 0.75 or so). Also, although including DBG-style search cost calculations to the BEFS-style model made little difference, using per-attempt search costs made the greedy/modest polymorphism even more robust.

The next step was to re-examine correlation. The modifications examined previously adjusted the correlation scheme in an attempt to make it as advantageous to the individual as possible. It seemed like a reasonable step to make correlation a completely individual-based characteristic, rather than strategy-based. Chapter V examined the results of this move, in part to determine reasonable upper limits on correlation values in the earlier models (DBG-style, etc.). An interesting result was that very few populations actually proceeded to the modest/greedy polymorphism when DBG-style correlation was an individual-based characteristic and initially random. The BEFS-style model also showed a tendency for the all-fair equilibrium to appear even when initially greatly

outnumbered, although the polymorphic equilibrium was by no means dominated.

Chapter V also acted as a transition of sorts into the group selection models discussed in Chapters VI and VII, since the individual-based correlation model can be interpreted as a group selection model. The simple group-selection models explored used multiple populations, each without any internal correlation. The idea was to see if the questions concerning correlation -- how high it needs to be, what scheme is most realistic, etc. -- could be bypassed using population structures as a substitute. The results were somewhat ambiguous. The trickle model did succeed in making any particular starting population proceed to the all-fair equilibrium, but only if the other populations had their starting strategy proportions randomly initialized. This move makes the global population (i.e. all individuals in all populations collectively) be in proportions which would naturally proceed to the all-fair equilibrium anyway if the populations were merged. This greatly lessens the usefulness of the trickle model, but it does show that a growing influx of fair individuals can make any starting population proceed to the all-fair state.

The haystack model also had some interesting but mixed results. The earlier switch away from the use of replicator dynamics allowed a given population to proceed to different final equilibria on different runs. The haystack model allowed differences in the probabilities of different final states to be exploited and increase the range of starting states which proceed to the all-fair equilibrium. If the probability of the all-fair equilibrium occurring is raised even slightly (say, by allowing a small amount of intra-population correlation, or via migration), then the haystack model might easily permit the all-fair equilibrium to become the end result for practically all initial proportion mixes.

Although my haystack implementation reduced the role of population growth often used in haystack models, using a population size growth model based on total population fitness nevertheless improved the performance of the all-fair equilibrium slightly. This improvement could probably be enhanced even further if the population growth rate were increased, or if the reproduction phase

were made slightly longer.

Does Skyrms' model demonstrate what it did originally, before the corrections of D'Arms and other authors? The best answer I can provide is: maybe. The wide and varied range of behaviour this deceptively simple model exhibits is staggering. Every modification I examined in this thesis brings out a new aspect of the model that was not initially apparent. The all-fair equilibrium loses its dominance under DBG-style modifications, regains it under BEFS at high correlation rates, only to lose it once again when cost calculations are introduced. Making correlation individual-based instead of strategy-based brings back the global dominance, but only for random initialization. Moving to multiple uncorrelated populations in a structure also brings back dominance, but only if migration of fair individuals passes some critical threshold. Haystack models can bring the all-fair dominance back, but only if reproduction phases allow significant growth differential, or alternately if fair individuals are given a slight boost within populations.

Every method which brings back the near global dominance of the all-fair equilibrium comes with some drawbacks, limitations, or innate assumptions which can be used to argue against that method. How convincing any of the methods I've examined are will depend greatly on one's intuitions concerning the plausibility of group selection, high correlation, or individual-based correlation. I certainly did not find a panacea, and while some might be convinced that Skyrms' model works as intended, others will disagree, or at least be hesitant in accepting the model. I personally fall into the latter category. As I changed the model to reflect new assumptions or paradigms, new behaviour seemed to spring from those changes which surprised me. I believe any model which accurately reflects human behaviour is going to involve a significant amount of complexity -- many such changes piled atop one another -- and so I must remain agnostic concerning the chances of the all-fair equilibrium in any such model.

Footnotes

¹ On occasion, figures may be duplicated further on to make comparisons between images simpler. In such cases, a more descriptive title is often used on the duplicate in order to highlight the differing characteristic; however, this list will only indicate the primary title.

² Technically, it is the parental investment for the two genders which is equalized rather than the genders themselves, but since the investment for both genders is often quite similar, a gender ratio close to 50/50 will be observed.

³ Actually, Fisher's argument on its own doesn't explain the absence of polymorphisms (e.g. half the population producing 9 females per male, while the other half produces 9 males per female, which would still result in a 50/50 split). However, the argument can easily be fine-tuned to explain this absence (Skyrms, 1996, p. 17).

⁴ I will often simply refer to 'the triangle' rather than using the more formal term 'state space.' I will also generally refer to 'individuals using the fair strategy' as just 'fair individuals,' or even simply as 'fairs.' I tend to be somewhat loose in my terminology at times, and a statement such as 'fair individuals would prefer to avoid greedies' simply means that fair individuals get a higher payoff playing against a strategy other than greedy. This is intended simply as a way of expressing things in a less wordy manner, and I do not intend to imply that individuals of any type necessarily have motivations, preferences, or desires. Skyrms' model can be applied to non-sentient creatures quite readily.

⁵ For values in between, the probabilities are adjusted as follows. Let e be the correlation constant, and suppose player 1 is using strategy S . If the original probability of player 2 using strategy S is $p(S)$, then the new probability (using the correlation factor) would be $p(S)+e \cdot p(\text{not-}S)$. Similarly, the probability of player 2 using the other strategies is reduced by a factor of e .

⁶ Since the model does not include mutation, there is no mechanism for introducing new types into the population. This means that the states which do not initially contain fair players (corresponding to the bottom line of the triangle) will still evolve to the modest/greedy polymorphism rather than the all-fair equilibrium.

⁷ I will be referring to the second (correlated) model simply as Skyrms' model. The earlier uncorrelated model is just a special case of the second.

⁸ Actually, like most computer simulations, it is strictly speaking deterministic. However, with a good pseudo-random number generator, results statistically similar to those involving true randomness can be obtained, and I will discuss them as if true randomness were involved.

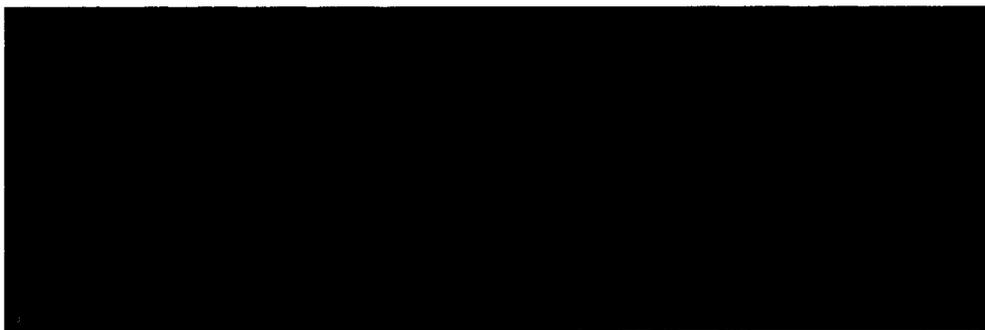
⁹ See Appendix 1 for details on population size, granularity, number of runs averaged, and other details for each image.

¹⁰ Also, some of this fuzziness in the separatrix is a result of using 'relaxed' stopping conditions to improve computational performance. See Appendix 1 for details on relaxed stopping conditions.

¹¹ The slight amount of non-blue at the bottom edge is due to extinction of all fair individuals, as discussed in section 4.2.1.

¹² Since the BEFS-style model with per-attempt cost had one of the most persistent basins of attraction for the modest/greedy polymorphism, I also did a simulation using random initial individual correlation for this type of model as well (cost factor 0.0444). However, analysis proved rather complex, and it detracted somewhat from the unity of the chapter, so I have elected to omit it from the thesis. However, since some rather unusual behaviour is seen in this model that is present nowhere else (such as the astonishing range over which the modest/greedy polymorphism is possible, or the development of all-modest populations), I include the result diagrams here without analysis for those curious.

BEFS-Style with Per-Attempt Cost (Factor 0.0444) and Random Initial Individual Correlation
Equilibria Fair Correlation (Cutoff 0.01) Greedy Correlation (Cutoff 0.01)



¹³ That Darwin was a proponent of group selection is contested by some authors (Ruse, 1989, ch2).

¹⁴ Actually, since some of the models considered here require persistent individuals, I will be using random sampling without replacement, but this will approximate Cohen-Eschel Case I quite closely if $\sum \text{final_population_sizes} \gg \text{initial_haystack_size}$.

¹⁵ Although just a single instance remained (out of 1326) at a 9% migration rate, so we might consider the polymorphism essentially eliminated at that point.

Bibliography

- Barrett, M., Eells, E., Fitelson, B., and Sober, E. (1999). "Models and Reality - A Review of Brian Skyrms's Evolution of the Social Contract," *Philosophy and Phenomenological Research* Vol. 59 No. 1: 237-241.
- Binmore, K.G. (1994). *Game Theory and the Social Contract, Volume 1: Playing Fair*. Cambridge: MIT Press.
- Cohen, D., and Eschel, I. (1976). "On the Founder Effect and the Evolution of Altruistic Traits," *Theoretical Population Biology* 10: 276-302.
- D'Arms, J., Batterman, R., and Gorny, K. (1998). "Game Theoretic Explanations and the Evolution of Justice," *Philosophy of Science* 65: 76-102.
- Danielson, P. (1992). *Artificial Morality: Virtuous Robots for Virtual Games*. London: Routledge.
- Dobzhansky, T. (1962). *Mankind Evolving: The Evolution of the Human Species*. New Haven: Yale University Press.
- Farber, P.L. (1994). *The Temptations of Evolutionary Ethics*. Berkeley: University of California Press.
- Hamilton, W.D. (1975). "Innate Social Aptitudes of Man: an Approach from Evolutionary Genetics," in Fox, R. (ed.), *Biosocial Anthropology* (133-153). London: Malaby Press.
- Maynard Smith, J. (1964). "Group Selection and Kin Selection," *Nature* 201: 1145-1146.
- Maynard Smith, J. (1998). "The Origin of Altruism," *Nature* 393: 639-640.
- Maynard Smith, J., and Price, G. (1973). "The Logic of Animal Conflicts," *Nature* 246: 15-18.
- Nagel, T. (1986). *The View from Nowhere*. New York: Oxford University Press.
- Nozick, R. (1993). *The Nature of Rationality*. Princeton: Princeton University Press.

- Richter, P.H., and Pietgen, H.-O. (1986). *The Beauty of Fractals*. Berlin: Springer-Verlag.
- Ruse, M. (1989). *The Darwinian Paradigm: Essays on its History, Philosophy, and Religious Implications*. New York: Routledge.
- Schroeder, D. (2005). "Evolutionary Ethics." *The Internet Encyclopedia of Philosophy*. Retrieved September 10, 2005, from <http://www.iep.utm.edu/e/evol-eth.htm>.
- Skyrms, B. (1996). *Evolution of the Social Contract*. New York: Cambridge University Press.
- Skyrms, B. (1998). "Mutual Aid: Darwin Meets The Logic of Decision," in Danielson, Peter (ed.), *Modeling Rationality, Morality, and Evolution* (379-407). New York: Oxford University Press.
- Skyrms, B. (1999). "Reply to Critics," *Philosophy and Phenomenological Research* Vol. 59 No. 1: 243-254.
- Skyrms, B. (2002). "Critical Commentary on Unto Others," *Philosophy and Phenomenological Research* Vol. 65 No. 3: 697-701.
- Sober, E., and Wilson, D.S. (1994). "A Critical Review of Philosophical Work on the Units of Selection Problem," *Philosophy of Science* Vol. 61 No. 4: 534-555.
- Sober, E., and Wilson, D.S. (1998). *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge: Harvard University Press.
- Sober, E., and Wilson, D.S. (2002). "Reply to Commentaries," *Philosophy and Phenomenological Research* Vol. 65 No. 3: 711-727.
- Taylor, P.D., and Jonker, L. B. (1978). "Evolutionarily Stable States and Game Dynamics," *Mathematical Biosciences* 40: 145-156.
- Tennant, N. (1999). "Sex and the Evolution of Fair-Dealing," *Philosophy of Science* Vol. 66 No. 3: 391-414.
- Woolcock, P.G. (1999). "The Case Against Evolutionary Ethics Today," in Maienschien, J., and Ruse, M. (eds.), *Biology and the Foundation of Ethics* (276-306). Cambridge: Cambridge University Press.

Appendix 1 - Details of Figures

This chart lists most of the relevant details for the evolutionary simulations that produced the figures in this thesis. See next page for column descriptions.

Fig.	Pop Size	Grn	Rep	Max Gen	R	CS	Correlation	Misc.
4.1	900	30	1	100	N	S	0	
4.2	900	30	10	100	N	S	0	
4.3	750	10	50	100	Y	S	0	
4.4	750	10	50	100	Y	S	0.2	
4.5	750	10	50	100	Y	D	<0, 0.2, -0.2>	
4.6	750	10	50	100	Y	D	<0, 0.2, -0.2>	DBG cost 0.3
4.7	750	10	50	100	Y	D	<0, 0.2, -0.2>	per-A cost 0.0444
4.8	750	10	50	100	Y	B	<0, -0.2, 0.2>	
4.9	750	10	50	100	Y	B	<0, -0.5, 0.5>	
4.10	750	10	50	100	Y	B	<0, -0.75, 0.75>	
4.11	750	10	50	100	Y	B	<0, -0.2, 0.2>	DBG cost 0.3
4.12	750	10	50	100	Y	B	<0, -0.75, 0.75>	DBG cost 0.3
4.13	750	10	50	100	Y	B	<0, -0.2, 0.2>	per-A cost 0.0444
4.14	750	10	50	100	Y	B	<0, -0.75, 0.75>	per-A cost 0.0444
4.15	750	10	50	100	Y	B	<0, -0.95, 0.95>	per-A cost 0.0444
5.2	750	10	50	100	Y	D	Individual	
5.3	750	10	50	100	Y	D	Individual	reused data for 5.2
5.4	750	10	50	100	Y	D	Individual	reused data for 5.2
5.5	750	10	50	100	Y	B	Individual	
5.6	750	10	50	100	Y	B	Individual	reused data for 5.5
5.7	750	10	50	100	Y	B	Individual	reused data for 5.5
5.8	750	10	50	100	Y	B	Individual	reused data for 5.5
7.1	500	10	10	50	N	S	0	trickle 0.01
7.2	500	10	10	50	N	S	0	trickle 0.02
7.3	500	10	10	50	N	S	0	trickle 0.03
7.4	50	1	10	50	N	S	0	haystack 10, 0.5
7.5	50	1	10	60	N	S	0	haystack 15, 0.5
7.6	50	1	10	100	N	S	0	haystack 25, 0.4
7.7	180	2	250	30	N	S	0	
7.8	180	2	250	30	N	S	0	reused data for 7.7
7.9	180	2	250	30	N	S	0.04	

Column Name	Description
Pop Size	The number of individuals in the population. In the case of haystack models, this indicates the initial population size.
Grn	The granularity of the image. Indicates the number of individuals differing between subsequent runs. Thus, (Pop Size/Grn)+1 yields the number of points calculated along each triangle side.
Rep	The number of separate runs averaged together for each point.
Max Gen	The maximum number of generations allowed to pass in a given run.
R	Indicated whether relaxed stopping conditions were used to improve calculation efficiency. When used, any given run is terminated before the maximum generation is reached if the mean individual fitness changes less than $(1/6) / (\text{Pop size})$ between iterations. This effectively terminates the run when the strategy proportions have changed by no more than a single individual.
CS	The correlation scheme used. S = Skyrms-style, D=DBG-style, B=BEFS-style
Correlation	The amount of internal correlation used. A value of 'Individual' indicates that individual-based correlation was used, randomly initialized between -1.0 to 1.0.
Misc	Indicates: (a) the type of cost calculations used, if any, and the cost coefficient (b) whether data from an earlier run was reused to create a subsequent figure (c) if a trickle model was used, the migration rate used. All trickle runs used 12 populations in a migration structure. (d) if a haystack model was used, the number of iterations between dispersal phases is shown, and the population delta. All haystack runs used 10 populations (haystacks) in a dispersal structure.