Stability and Evolution in Rock-paper-scissors Ecologies

Richard Mansfield

Submitted in fulfilment of the requirements for the degree of Doctor of Philosophy in Computer Science.

Victoria University of Wellington 2006

Abstract

Four aspects of rock-paper-scissors ecosystems are considered. In the first, the effects of variations in the population's spatial structure are described. With a lattice-based spatial structure, three-species coexistence is unstable when dispersal is long-range, but becomes stable at a critical threshold as the dispersal distance is gradually reduced. A continuous-space model is constructed and is shown to reproduce the three-species stability of the lattice model when interactions are local. But unlike the lattice model, three-species coexistence can be stable in the continuous space model even when dispersal is long-range.

In the second investigation, the pair approximation technique is applied to the rock-paper-scissors system. The resulting equations fail to predict the stability of the three-species state. A local structure approximation based on four neighbouring sites is shown to produce a more accurate result.

The third contribution is the presentation of two models in which rock-paperscissors is able to evolve from a simple two-species system. In both cases, a two-species competitive cellular automaton is augmented with individual variation in a phenotypic trait. In the first model, the trait determines the individual's investment in interspecific versus intraspecific competition, and in the second model, the trait determines the individual's investment in the production of an interspecific toxin versus its growth rate. When interactions are local, selection can cause a divergence in the values of both traits such that only individuals with extreme levels of the trait survive, while those with intermediate levels of the trait die out. After this divergence, the resulting polymorphic community competes in an intransitive rock-paper-scissors cycle.

Finally, an assessment is made of the claim that competitive restraint in rockpaper-scissors systems is caused by selection for stable subcommunities in a process of community-level selection. It is argued that members of unstable subcommunities may be as fit or fitter than identical individuals in stable subcommunities, and that the invocation of the higher-level selective force does not add to the understanding of the process of competitive restraint. Measurements of community stability in cellular automata models of rock-paperscissors are undertaken and these provide no evidence for a community-level selective force promoting restraint.

Acknowledgements

First I'd like to thank Marcus Frean for providing guidance and useful advice, and for creating an excellent environment in which to be a research student.

Thanks also to Ed Abraham for reading the thesis at the last minute and providing useful comments.

Thanks to all the people I've shared offices with: Gareth Baxter, Phillip Boyle, Jeromé Dolman, Tim Field, Sergio Hernández, Ian Liang, Mark Rickerby, Russell Tod, and Ryan Woodard, for entertaining discussions and distractions. Thanks also to all the regular attendees of the Festival of Doubt (Wellington Chapter), who gave stimulating talks and listened to a couple of off-topic ones from me; thanks to Vivian McPhail, for all the arguments about consciousness over lunch; and to my flatmates Andy, James, Russell and Vanessa for creating a happy domestic environment.

Thanks to the School of Mathematics, Statistics, and Computer Science for the use of the grid, which was essential for much of the experimental data, and to the school's programmers, because the systems always seemed to work.

Finally, the thesis would not have been possible without the assistance of a Victoria University postgraduate scholarship.

Contents

1	Roc	ock-paper-scissors ecosystems		
	1.1	The mean field model		
	1.2	Stocha	astic cellular automaton model	4
	1.3	Finite	-population model with long-range dispersal	6
	1.4	Outlir	ne of the thesis	7
Ι	Th	e effe	cts of spatial structure in agent-based models	13
2	RPS	in con	tinuous space	15
	2.1	A con	tinuous-space model for RPS	16
	2.2	Comp	parison with the cellular automaton model	19
	2.3	Long-	range dispersal	22
		2.3.1	Continuous-space model with long-range dispersal	22
		2.3.2	1-1 versus 1-n interactions	23
		2.3.3	Fixed populations and equal interaction probabilities	26
	2.4	Discu	ssion	27
3	Inte	ermedia	ite-range dispersal	29
	3.1	Time f	to extinction (TTE) experiments	30
	3.2	Increa	sing the grid size	31
		3.2.1	Two reasons why large grids increase stability	31

		3.2.2	RPS versus random drift	32
		3.2.3	Invasion rates and averaging of TTE	33
		3.2.4	TTE increases with grid size	33
		3.2.5	Trends in average TTE	33
		3.2.6	Distribution of TTE	34
	3.3	Decre	asing the dispersal range	35
		3.3.1	The effect of neighbourhood size on TTE	37
		3.3.2	Neighbourhood size and TTE distributions	40
	3.4	Comb	inations of grid size and dispersal distance	41
		3.4.1	Relative importance of grid size and neighbourhood	41
		3.4.2	Effect of invasion rates and cluster size	43
		3.4.3	Stability is maximised for small, but not minimum neigh-	12
			bournoods	40
	3.5	Summ	nary and discussion	43 44
	3.5	Summ	nary and discussion	43
II	3.5 A]	Summ pprox	imations of spatial structure	43 44 47
II 4	3.5 Aj Pair	Summ pprox	imations for RPS ecosystems	43 44 47 49
II 4	3.5 A J Pair 4.1	Summ pprox approx Pair a	bournoods	 43 44 47 49 50
II 4	 3.5 A] Pair 4.1 4.2 	Summ pproxi approx Pair a Appli	bournoods	43 44 47 49 50 51
II 4	 3.5 A] Pair 4.1 4.2 	Summ pprox approx Pair a Appli 4.2.1	bournoods	43 44 47 49 50 51 51
II 4	 3.5 A] Pair 4.1 4.2 	Sumn pprox approx Pair a Appli 4.2.1 4.2.2	bournoods bournoods hary and discussion bournoods imations of spatial structure imations for RPS ecosystems nd local structure approximations cation to RPS Pair transformation dynamics Symmetry in pair types	43 44 47 49 50 51 51 51
11 4	 3.5 A] Pair 4.1 4.2 	Summ pproxit approxit Pair at Appli 4.2.1 4.2.2 4.2.3	bournoods	43 44 47 49 50 51 51 51 51 52
II 4	 3.5 A] Pair 4.1 4.2 	Summ pprox approx Pair a Appli 4.2.1 4.2.2 4.2.3 4.2.4	bournoods	43 44 47 49 50 51 51 51 51 52 53
II 4	 3.5 A] Pair 4.1 4.2 4.3 	Summ pproxi approxi Pair a Appli 4.2.1 4.2.2 4.2.3 4.2.4 Deriva	bournoods	43 44 47 49 50 51 51 51 51 51 52 53 54
II 4	 3.5 A] Pair 4.1 4.2 4.3 	Summ pprox approx Pair a Appli 4.2.1 4.2.2 4.2.3 4.2.4 Deriva 4.3.1	bournoods	43 44 47 49 50 51 51 51 51 51 52 53 54 54
II 4	 3.5 A] Pair 4.1 4.2 4.3 	Summ pprox: approx: Pair a: Appli 4.2.1 4.2.2 4.2.3 4.2.4 Derive 4.3.1 4.3.2	bournoods	43 44 47 49 50 51 51 51 51 51 52 53 54 54 55
II 4	 3.5 A] Pair 4.1 4.2 4.3 	Summ pproxi approxi Pair a Appli 4.2.1 4.2.2 4.2.3 4.2.4 Deriva 4.3.1 4.3.2 4.3.3	bournoods bournoods hary and discussion bournoods imations of spatial structure imations for RPS ecosystems nd local structure approximations cation to RPS Pair transformation dynamics Symmetry in pair types Counting of pairs Homogeneous and heterogeneous pairs ation of pair correlation equations Rate of change of homogeneous pairs Rate of change of heterogeneous pairs Rates of transformation of pairs	43 44 47 49 50 51 51 51 51 51 51 52 53 54 54 55 55

		4.3.4	The pair closure	56
		4.3.5	Rate of change of singleton densities	58
	4.4	Failur	e of the pair approximation for RPS	59
		4.4.1	Importance of neighbourhood size	60
		4.4.2	Extinctions more common than coexistence	60
		4.4.3	Effects of initial pair densities	62
		4.4.4	Pair approximation compared to SCA simulation	63
		4.4.5	Summary	63
	4.5	2×2	approximations	64
		4.5.1	RPS and 2 \times 2 blocks	65
		4.5.2	Block densities, ordered and unordered	66
		4.5.3	Rate of change of block types	67
		4.5.4	Block transformation rates	68
		4.5.5	The '2 \times 2 block' approximation $\hfill\hfilt$	70
	4.6	Behav	riour of the 2 \times 2 model \ldots	71
		4.6.1	Stability	72
		4.6.2	Ecosystem composition	72
		4.6.3	Neighbour correlations and species clustering	74
	4.7	Discu	ssion	76
5	Dire	ect spat	ial stability approximation	79
	5.1	Explic	it spatial stability	80
	5.2	Behav	iour of the approximation	80
		5.2.1	Extinction risk	81
		5.2.2	Predictions of extinction risk	82

II	ΙΤ	he ev	olution of competitive cycles	85
6	Intr	aspecif	ic competition	87
	6.1	Mode	lling speciation	88
		6.1.1	Sympatric speciation	88
		6.1.2	Assumptions of adaptive dynamics models	89
		6.1.3	Evolution in AD models	91
		6.1.4	Adaptive dynamics and RPS systems	91
	6.2	The ir	ntraspecific competition model	92
		6.2.1	State description	92
		6.2.2	Ecological interactions	93
		6.2.3	The evolutionary process	94
		6.2.4	Initial state description	94
	6.3	Behav	viour of the model	95
		6.3.1	Stable coexistence	95
		6.3.2	Early oscillations in species densities	95
		6.3.3	Evolutionary branching	98
		6.3.4	Attracting states in the model	102
	6.4	Realis	stic scenarios	107
		6.4.1	Scenario types	107
		6.4.2	Simulations of scenarios 1 and 2	108
		6.4.3	Plausibility of mutation rates	112
	6.5	The ir	ntraspecific competitiveness trait	113
	6.6	Sumn	nary	115
7	Den	sity sta	ability in the intra model	117
	7.1	A non	nspatial intraspecific competition model	118
		7.1.1	Model definition	118
		7.1.2	Ecological dynamics	119

		7.1.3	Evolutionary dynamics	. 120
		7.1.4	Evolutionary and ecological timescales	. 121
		7.1.5	Summary of model assumptions	. 121
	7.2	Proper	rties of the nonspatial model	. 122
	7.3	Nonsp	patial finite-population simulations	. 124
		7.3.1	Intraspecific competition vs. RPS	. 124
		7.3.2	Mean traits versus trait distribution	. 126
		7.3.3	Flat trait distributions in nonspatial simulations	. 127
	7.4	Impor	tance of spatial structure for coexistence	. 129
	7.5	Summ	ary	. 131
0	Dia		colocition in the intro model	100
8	DISI	ruptive	selection in the intra model	133
	8.1	Spatia	l interactions create local clustering	. 134
	8.2	The th	ree-cell model	. 134
		8.2.1	A measure of short-term invasion fitness	. 134
		8.2.2	Quantification of clustering	. 135
		8.2.3	Neighbour relationships	. 136
		8.2.4	Possible phenotypes in the initial state	. 136
		8.2.5	Nine initial states	. 136
		8.2.6	Approximation of the system dynamics	. 137
		8.2.7	Interaction fitness and invasion probabilities	. 138
		8.2.8	Expected two-interaction fitness	. 139
		8.2.9	Total expected short-term invasion fitness	. 140
		8.2.10	Invasion fitness in clustered environments	. 140
	8.3	The fa	stest invaders survive	. 140
	8.4	Asym	metric competition	. 142
	8.5	Summ	ary	. 142
9	Cyc	lic com	petition among toxin-producers	145

9.1	Mode	lling toxin-based competition
	9.1.1	Theoretical models of microbial toxin systems 146
	9.1.2	Advantages of the intraspecific model
	9.1.3	Application of the intraspecific model
9.2	A two	-species toxin-growth tradeoff model
9.3	Specie	es densities
	9.3.1	Coexistence and stability
	9.3.2	Stability in spatially unstructured populations 154
	9.3.3	Net invasion rates for pure strategies
	9.3.4	Mean-field model
	9.3.5	Local clustering of species
9.4	Distril	oution of toxin-producers
	9.4.1	Classification of toxin-production distributions 164
	9.4.2	Saturation of toxin production for one species 168
	9.4.3	Evolutionary branching
	9.4.4	Summary of equilibrium behaviour
9.5	Invasi	on of an established population
9.6	Gener	alised toxin-growth model
	9.6.1	Mean field equations
	9.6.2	Species densities in the general model
	9.6.3	Distribution of toxin-producers in the general model 176
	9.6.4	Invasion of an established population
9.7	Summ	nary

IV Competitive restraint

183

10	0 Restraint and community level selection					
	10.1 Levels of selection	186				

	10.1.1	Group selection	186
	10.1.2	Group selection and altruism	187
	10.1.3	Multilevel selection theory	187
	10.1.4	Group selection or individual selection?	188
	10.1.5	From group selection to community selection	188
10.2	Popula	ation structure	189
	10.2.1	What is a group?	189
	10.2.2	What is a (sub)community?	189
10.3	Spatia	lly continuous landscapes	190
	10.3.1	The evolution of altruism in viscous populations	191
	10.3.2	Fluctuations in population size	191
	10.3.3	Altruism and competitive restraint	192
	10.3.4	Provisional examples of selection in continuous com- munities	192
10.4	Comm	nunities in cellular automata	193
	10.4.1	The community selection process	194
	10.4.2	Community definition using interactions and traits	194
	10.4.3	Community definition using spatial self-structuring	196
	10.4.4	Subcommunity interaction and heritability	197
10.5	Conclu	usion	197
11 Rest	raint a	nd stability	199
11.1	Time t	o extinction on small grids	201
	11.1.1	Invasion rates and time to extinction	202
	11.1.2	Individual and subcommunity selection	203
	11.1.3	Variation in extinction times at the restraint equilibrium	205
11.2	Minim	num species density	205
11.3	Variati	ion in species densities	208
		1	

11.4 Cluster size	. 210
11.5 Summary	. 212
12 Conclusions	215
Appendix: The adaptive dynamics method	221
Bibliography	223

Chapter 1

Rock-paper-scissors ecosystems

Rock-paper-scissors (RPS) is a children's game in which two players each make one of three moves: rock, paper, or scissors. Rock wins against scissors, scissors wins against paper, and paper wins against rock; if both players make the same move, it is a draw. In game theory, a simple version of RPS can be described by the following matrix, describing the payoff to player 1:

	Player 2			
Player 1	r	S	р	
r	0	1	-1	
S	-1	0	1	
р	1	-1	0	

Player 1 gains a point when playing rock (r) against scissors (s), but loses a point playing rock against paper (p). RPS is the simplest game in which the winner is decided by an intransitive dominance relationship between the game's moves. When the game is played repeatedly, none of the three pure strategies r, s, or p has an advantage. If we include mixed strategies, where each of the moves can be played with some probability, then the strategy in which the three moves are played with equal probability cannot be beaten.

In evolutionary game theory, members of a population play pure strategies against each other, and reproduce in proportion to their relative success. Strategies are interpreted as different phenotypes or species and the underlying gene frequencies in the population are ignored. In evolutionary rock-paperscissors, there is no evolutionary stable strategy; a population made up of any one phenotype can always be invaded by one of the others (Maynard Smith, 1982). Depending on the exact payoff structure, the intransitivity in the game can lead to either coexistence of all three strategies in constant proportions, or to endless oscillations in their populations (May and Leonard, 1975; Gilpin, 1975b).

The oscillations in the three populations are a direct result of the intransitivity, because an increase in the number of rock-playing individuals comes at the expense of scissors-players. Fewer scissors leads to more paper, which in turn leads to a decrease in rocks. So even though rock-players compete with scissors-players, the rock-players benefit from the presence of scissors because of the effect that scissors have on paper.

In natural ecosystems, indirect effects like this are common. Real multispecies communities tend to contain multiple intransitive loops that transform apparently competitive relationships into facilitative ones. One study has calculated that in large competitive networks, 20-40% of interactions that would be detrimental, if the two species involved were taken in isolation, turn out to be beneficial when indirect effects are taken into account (Stone and Roberts, 1991).

It is hoped that the study of the simple RPS ecosystem will lead to a better understanding of the counterintuitive consequences of intransitive species relationships in general. But this is not the only motivation, because several examples of simple intransitive competition among three phenotypes are known to exist in nature. The most well-known involves three mating strategies of the male side-blotched lizard, in which the frequency of each type in the overall population fluctuates as predicted by the evolutionary models (Sinervo and Lively, 1996). Three-species intransitive competition has also been observed among corals and other sessile (non-mobile) animals that compete for space in reef environments (Buss and Jackson, 1979; Buss, 1980; Johnson, 1997), and among strains of bacteria which produce toxins called colicins (Kerr et al., 2002; Kirkup and Riley, 2004).

Spatial interactions in ecological models

Traditional models based on evolutionary game theory assume a population that mixes completely at random, in such a way that any individual has an equal probability of meeting any other individual. While this assumption simplifies the models, accuracy may be sacrificed in the process, because spatial structure can promote diversity and stability both in ecological models (Hassell et al., 1994; Durrett and Levin, 1994a), and when added to traditional game theory, even in games such as the prisoner's dilemma where one strategy dominates in the non-spatial form (Nowak and May, 1992).

1.1. THE MEAN FIELD MODEL

In RPS, whether the population mixes randomly or is spatially structured has a crucial effect on the survival of the ecosystem, because spatially structured populations tend not to oscillate, but remain stable over long periods of time, with nearly constant proportions of the three species.

These spatial effects are illustrated by Frean and Abraham (2001) in three basic models which are used throughout this thesis and which are summarised in the next three sections.

1.1 The mean field model

Under the 'mean-field' assumption, the population is completely unstructured so that at any time, an individual is equally likely to interact with any other individual.

Application of the mean field model to the rock-paper-scissors ecology is covered by Frean and Abraham (2001). The actual numbers of rock, scissors, and paper individuals are abstracted away and the model tracks their 'densities', the number of individuals as proportion of the total population. The model consists of three differential equations describing the rates of change of the three species densities, which can be derived from what Hofbauer and Sigmund (1998) call the 'replicator dynamics'. For example, the density of rocks, ρ_r , varies according to the following equation:

$$\frac{\mathrm{d}\rho_{\rm r}}{\mathrm{d}t} = \rho_{\rm r}\rho_{\rm s}r_{\rm r} - \rho_{\rm p}\rho_{\rm r}r_{\rm p} \tag{1.1}$$

where r_r is the rate at which rocks grow over scissors, and r_p is the rate at which paper grows over rocks. The first term, $\rho_r \rho_s r_r$, represents the increase in the proportion of rocks in the population from interactions in which rocks invade space previously occupied by scissors, and the second term, $-\rho_p \rho_r r_p$ is the decrease in the proportion of rocks resulting from the invasions by paper of space previously occupied by rocks. Similar equations can be written for the dynamics of the densities of the other two species, ρ_s and ρ_p .

The mean field model predicts that the densities travel in stable orbits around a fixed point as in figure 1.1. At this fixed point, the three densities are

$$\rho_{\rm r} = \frac{r_{\rm s}}{r_{\rm r} + r_{\rm s} + r_{\rm p}}, \qquad \rho_{\rm s} = \frac{r_{\rm p}}{r_{\rm r} + r_{\rm s} + r_{\rm p}}, \qquad \rho_{\rm p} = \frac{r_{\rm r}}{r_{\rm r} + r_{\rm s} + r_{\rm p}}$$

Many real ecosystems do not conform with the mean-field assumption. They have spatial structure and individuals are typically limited in the distance that



Figure 1.1: Oscillations in species densities under the mean-field RPS model of equation (1.1), for the invasion rates $r_r = 0.4$, $r_s = 0.6$, $r_p = 1$. The five small circles show five initial species densities (ρ_r , ρ_s , ρ_p) from the set {(0.3,0.5,0.2), (0.35,0.4,0.25), (0.4,0.3,0.3), (0.45,0.2,0.35), (0.5,0.1,0.4)} and the lines show the trajectories followed by the species densities (in a clockwise direction) from these initial points. The leftmost circle is the fixed point of the equations from which the densities remain constant.

they can move, so in practice, the probability that two individuals *i* and *j* will interact at a particular time will vary with *i* and *j*. Under these circumstances we should expect the mean field model to be limited in its ability to predict the long-term behaviour of the ecosystem.

1.2 Stochastic cellular automaton model

Stochastic cellular automata (Silvertown et al., 1992; Durrett and Levin, 1994b) are often used in ecological modelling in order to overcome the limitations of the mean field approach when dealing with spatially structured populations. SCAs are agent-based: they explicitly represent every individual in the population by locating them at the vertices of a lattice or graph. Only individuals that are joined by an edge are able to interact directly; the resulting localisation of interactions is what captures the spatial structure of the population.

The description that follows is similar to that of the SCA model given by Frean and Abraham (2001). *N* individuals, each of which has a species in $\{r, s, p\}$, are located at the vertices of a square grid with *N* sites. The boundaries of the



Figure 1.2: Comparison of the mean field prediction with a spatially explicit simulation, on a cellular automaton, of the RPS system using invasion rates $r_r = 0.4$, $r_s = 0.6$, $r_p = 1$. The initial population densities (0.45,0.2,0.35) are marked with an open circle, the mean field trajectory is shown by the dotted line, the trajectory of the SCA simulation is shown by the solid line, and the mean field fixed point is marked with a star. Under the SCA model, the densities are attracted to stable proportions very close to those at the fixed point of the mean field equations.

grid are periodic, so that sites on the top edge of the grid are adjacent to those at the bottom, and sites at the left edge are adjacent to those at the right.

At each timestep, two individuals are chosen to interact. One site is chosen uniformly from the entire population and a second is chosen from among the direct neighbours of the first. The occupant of the first site replicates into the second with a probability determined by the species of the two occupants and the invasion rates: r invades s with probability r_r , s invades p with probability r_s , and p invades r with probability r_p ; for any other combination of species in the two sites there is no invasion and the grid state is unchanged.

A spatial simulation can be compared to the mean field model by initialising the sites of the SCA so that the three species are in the same proportion as the initial densities in the mean field equations (figure 1.2). The mean field trajectory follows a stable orbit, but the trajectory of the SCA settles into an attractor near to the mean field fixed point. In RPS ecosystems, spatial structure has a stabilising influence on the composition of the population.



Figure 1.3: Simulation of 1000 randomly-mixing individuals with invasion rates $r_r = 0.4$, $r_s = 0.6$, $r_p = 1$. The initial population densities are marked by the open circle, near the fixed point of the mean field model. The densities go through increasingly large oscillations in which one species dominates until eventually one species (in this case paper) goes extinct, after which its prey species takes over and the densities settle at the rest point marked by the black circle.

1.3 Finite-population model with long-range dispersal

The mean-field model is effectively an infinite-population model, because there is no limit to how small the proportions described by ρ_r , ρ_s , and ρ_p can become, provided they are greater than zero to begin with. The spatial SCA, on the other hand, has a finite population of size *N*, so species densities are always multiples of 1/N.

If we take the finite population of the SCA and the random-mixing assumption of the mean field, we get a third model in which a pool of individuals compete for slices of their fixed-size population. Individuals are said to have 'long-range dispersal' because rather than replicating into neighbouring sites, they can replicate into any other site, as if they could disperse into far-flung areas on a landscape.

In this case, oscillations in the densities slowly increase in amplitude until one species goes extinct, followed soon after by a second species (figure 1.3). The species with the fastest invasion rate usually goes extinct, and the 'winning' species is usually the slowest. Frean and Abraham (2001) describe the phenomenon as 'the survival of the weakest'.

1.4 Outline of the thesis

The thesis is in four parts, all of which deal with RPS, and in some cases other ecosystems with intransitive competition. Part I presents some more realistic models of spatial structure than those described in sections 1.2 and 1.3. Part II describes some efficient methods for approximating spatial structure without explicit representation of every agent in the system, and applies these methods to RPS systems. Part III introduces models for the evolution of RPS ecosystems through speciation events, and Part IV is a partial critique of a particular explanation of apparently altruistic behaviour in RPS. The major contributions of the thesis are described, part by part, in the remainder of this section.

Part I: Spatially-structured populations

Completely well-mixed natural populations are uncommon, so the mean field approximation often leads to inaccuracies by underestimating the degree of spatial structure. However, lattice-based approximations such as the SCA can overestimate the amount of spatial structure because it is also rare for individuals to interact with a fixed set of neighbours. Chapter 2 attempts to gauge the accuracy of the lattice approach. An agent-based, continuous-space, nonlattice-based model of RPS is introduced and compared to the SCA model given in section 1.2 and the non-spatial model of section 1.3, and concludes that

- (1) When dispersal is localised, the results of the grid-based model approximate the non-lattice model quite well;
- (2) The finite-population, random-mixing model of section 1.3 approximates long-range dispersal in continuous-space quite poorly, at least for nonmobile organisms.

Both the well-mixed and local grid-based SCA models suffer from a lack of realism in their treatment of space, and there is also a huge disparity between the likelihood of ecosystem collapse under the two models when populations are finite. For the purely local SCA, stable coexistence (as in figure 1.2) is almost certain, and for the long-range, finite-population model, extinctions

(as in figure 1.3) are almost certain. In chapter 3, the effect of variation in neighbourhood size on ecosystem stability is assessed using simulation, to establish the nature of the transition between these two outcomes. The major results are:

- (3) When dispersal is changed slowly from long-range to local, there is a point at which the stability of the ecosystem increases at a greater than exponential rate;
- (4) Very short-range dispersal produces more stability than completely local dispersal.

Part II: Approximating spatial structure in non-agent-based models

Agent-based models (like those in sections 1.2 and 1.3) represent every individual and simulate every interaction, which is time-consuming because significant computation is required. The mean field model is more efficient because it models density changes directly, abstracting away from the details of individual interactions. Faster simulation makes for better generalisation and understanding, because a thorough exploration of the model space becomes possible.

The mean field model described in section 1.1 has an unrealistic treatment of space, and as a consequence it fails to predict the robust stability in a RPS ecosystem. Ecological pair approximation models (Matsuda et al., 1987) have shown that it is possible to model aspects of a spatially-structured population without explicit agent representation, by describing the dynamics of pairs of neighbouring individuals. In chapter 4, I apply the method to the RPS ecosystem, but find that

- (5) An approximation based on neighbouring pairs fails to predict the stability of the RPS system; it usually predicts extinctions instead; and
- (6) A more detailed approximation based on a larger group of neighbours predicts stability, but gives incorrect species densities at equilibrium.

A more ad-hoc model that forces stability on the mean field model is introduced in chapter 5, with the result that

(7) It is easy to add stability to an efficient mean-field type model of RPS, but it is not easy to reproduce the size of the early oscillations in the full

1.4. OUTLINE OF THE THESIS

spatial RPS model.

Part III: Evolution of cyclic species networks

In the well-known prisoner's dilemma game (Axelrod, 1984), defectors do better than cooperators when the two strategies are in direct competition, and in the well-mixed evolutionary version of the game, cooperators are driven to extinction. Some attempts to avoid these tragic consequences have involved adding a third strategy in addition to cooperators and defectors, which can transform the game into a nontransitive competition network in which diversity is preserved. For example, Szabó and Hauert (2002) investigated the addition of 'loners', who choose not to play the game, and found that this can create a cyclic network in which defectors beat cooperators, cooperators beat loners, and loners beat defectors. In another example, the 'policing' behaviour modelled in Frank (1995), by which social insects suppress the effects of competition, can be added to the standard prisoner's dilemma resulting in a three-strategy game in which defectors can beat cooperators, cooperators can beat 'police', and police can beat defectors (Mansfield, 2001).

Part III introduces two models, both of which begin in a state with two competing, asymmetric species. In the absence of evolution, one species is stronger and drives the other to extinction just like in the standard prisoner's dilemma. With the addition of a trait that is subject to natural selection, however, a speciation event may occur, and transform the system into a stable polymorphic intransitive network.

In the first model, a trait is introduced for intraspecific competition. Individuals vary in the amount of effort they expend competing against members of their own species versus the amount of effort spent competing against members of the other species. Results show that

- (8) A two-species ecosystem with evolving intraspecific competitiveness often results in a robustly stable community, even when the two species are mismatched in their inherent strengths;
- (9) When the population is spatially structured, a branching event in the trait of at least one species splits individuals of the species into two morphs; and
- (10) The two morphs of the polymorphic species, along with the other species, compete in an intransitive cycle.

Few if any species directly trade interspecific competitiveness against intraspecific competitiveness in this way. But toxin-producing species do something similar; they trade toxin production against growth rate. The second, more realistic model described in Part III is a model of a two-species system which is augmented with an individual-level trait for the production of an interspecific toxin. The following conclusions are drawn:

- (11) A two-species system with evolving interspecific toxin production is stable for a wide range of toxin strengths and growth rates.
- (12) One species' toxin usually evolves to a minimum or maximum, and the other species frequently undergoes a branching event and becomes polymorphic, in which case competition among morphs is intransitive.
- (13) The model suggests a possible mechanism for the origin of the high diversity which is observed in most natural communities in which toxin-production is widespread.

Part IV: Competitive restraint

In all three of the RPS models, the fastest invader usually ends up with the smallest population. Consequently, a species does better by slowing down its rate of invasion; it increases its density and decreases its risk of extinction. Frean and Abraham (2001) ran the RPS system while allowing rock's invasion rate to evolve while keeping the other two invasion rates fixed, and found that the density of rocks decreases as their average invasion rate increases.

Johnson and Seinen (2002) found a more interesting result in a spatial RPS system similar to that of section 1.2. When one of the three invasion rates is allowed to evolve while the other two are held fixed, the evolving rate does not increase forever but reaches a maximum value. They describe this phenomenon as 'competitive restraint', because the evolving species reaches a point at which it has to restrain its competitiveness, or else it goes extinct due to the 'survival of the weakest' property of the system. The evolving species is not, of course, 'voluntarily' restraining competitiveness, but individuals that become very competitive tend to go extinct in their local region, while those that remain slightly less competitive survive.

Johnson and Seinen take this explanation up to a higher level and claim that the occurrence of competitive restraint is caused by competition among multispecies subcommunities, and that in the competition among subcommunities, those that are the most stable tend to win. They claim that evolving-species in-

1.4. OUTLINE OF THE THESIS

dividuals face two opposing selective forces: individual-level selection, which tends to increase their competitiveness, and community-level selection, which tends to decrease competitiveness in order to make the community more stable.

I examine this claim in Part IV, and argue the following points:

- (14) The invocation of a higher-level selective force to explain the occurrence of competitive restraint is unnecessary and does not add to our understanding of the process.
- (15) Measurement of community stability (using different definitions) provides no evidence for the idea that competitive restraint is caused by a process in which community and individual-level selective forces are in balance at the equilibrium.

Part I

The effects of spatial structure in agent-based models

Chapter 2

Rock-paper-scissors in continuous space

Agent-based simulations of rock-paper-scissors ecologies in the published literature have typically been performed using stochastic cellular automata. For example, SCA models are used by Johnson (1997), Durrett and Levin (1998), Frean and Abraham (2001), Johnson and Seinen (2002), and Szabó et al. (2004). These authors have used lattice-based models under the assumption that they are good approximations to real ecosystems. But in real ecosystems, although individual organisms usually interact within spatial constraints, they are most frequently not restricted to interactions with a fixed set of neighbours in quite the same way as agents on a lattice.

The lattice assumption is useful because it provides spatial restrictions without the computational overhead of keeping track of every agent's precise location, and in many cases, the spatial information in the lattice models is sufficient to reproduce an ecosystem's essential behaviour. For example, results of the RPS laboratory experiments of Kerr et al. (2002) are predicted well by the model in section 1.2 on page 4.

Because of the speed of SCA models, I have used them throughout this thesis. But in order to confirm that the SCA is a reasonable approximation to the biological reality, in this chapter I present a continuous-space model that is capable of reproducing most of the essential features of RPS ecologies.

While continuous-space models are computationally expensive, they are being used more often in ecological modelling (for example, see Donalson and Nisbet (1999)) as the available computing power has increased. The model presented in section 2.1 is a simple model in which individuals do not move,



Figure 2.1: Probability distribution for the random distance z between parent and child.

but in which offspring are born, and interact, at arbitrary locations in space.

A comparison of this model with a SCA model, described in sections 2.2 and 2.3, shows that the two models produce very similar results in the case of local dispersal, but different results for long-range dispersal. Some implications of this discrepancy are discussed in section 2.4.

2.1 A continuous-space model for RPS

Each individual in the system is described completely by

- (1) a point location in two-dimensional space, and
- (2) a species in $\{r, s, p\}$.

Individuals have sizes described by a circle with identical radii around their point location. No two individuals' circles are allowed to overlap.

In the SCA models, the grid size is defined in terms of the number of individual sites it contains. In the continuous-space model I follow a similar convention and give individuals' circles a fixed diameter of one, so that it is possible to describe the total size of the world as a multiple of the size of a single individual. The space itself, like in the SCA, is defined using a square with periodic boundary conditions, which represents the surface of a torus.

The model is initialised by scattering a number of individuals at random across the space so that none of their circles overlap. The fixed invasion rates of the three species, (r_r, r_s, r_p) , are then used to update the world, by repeating the following three steps:

2.1. A CONTINUOUS-SPACE MODEL FOR RPS

- (1) A single parent individual is chosen randomly, with a uniform distribution, from all the individuals currently in the system.
- (2) A child is created, with the same species as the parent chosen in step 1, and this child is propagated in a random direction from the parent out to a distance *z* from the parent's centre, where *z* is chosen from an exponential distribution with a minimum of 1 and a mean of 2 as shown in figure 2.1. Because of the minimum distance, the child will never overlap with its parent, and the shape of the distribution ensures that interactions remain local.
- (3) The fate of the child is determined by its invasion rate *r* versus its prey species, and the locations of the existing individuals in the system as follows:
 - (a) If no existing individuals overlap the child's radius, then the child is successfully born and remains at the location chosen in step 2.
 - (b) If any of the individuals that overlap the child are of the same species as the child, or of the child's predator species, then the birth event is unsuccessful, the child is removed and the state of the world is unchanged.
 - (c) If all *n* overlapping individuals are members of the child's prey species, then with probability r^n , the child is successfully born and the *n* overlapping individuals are removed, otherwise the birth event is unsuccessful and the state of the world is unchanged.

Each iteration through steps (1)-(3) will be called an interaction, and the system time is measured in the number of interactions since the initialisation of the system.*

Figure 2.2 shows some snapshots of the state of the world taken from a simulation initialised with a small population. The distribution of the distance between parent and child means that the space begins filling up only in areas close to the initial individuals.

^{*}In many ecological SCA models, system time is defined by grouping a number of interactions equal to the total population size, like the 'epochs' used in Frean and Abraham (2001), or the synchronous timesteps in Durrett and Levin (1994b). In the following chapters I use SCA 'generations' defined in terms of the population size, but for the remainder of this chapter I use single interactions to measure time even when referring to SCA models.



Figure 2.2: Typical behaviour of the continuous space model over time. The figures show snapshots of the world at six stages in a simulation with an initial population of 250 individuals, a 150×150 space, and invasion rates (1,0.6,0.4). Red, green and blue dots mark locations of rock, scissors, and paper individuals.



Figure 2.3: Initial and long-term fluctuations in species densities for the continuousspace model (top row, 150 × 150 space) and CA (bottom row, 115 × 115 grid) with invasion rates (1, 0.6, 0.4). For the continuous-space model, the black line shows the size of the total population, and the red, green, and blue lines show the rock, scissors and paper population sizes. In the CA grid, the total population size is fixed at 13,225 (115 × 115) so that it is close to the maximum population size in the continuous-space simulation.

2.2 Comparison with the cellular automaton model

When the continuous space has filled up, the fastest invader has an initial population advantage, which is why there are large red areas (occupied by rock) in the fourth and fifth diagrams in figure 2.2. Eventually the slowest invader (in this case paper) tends to take over, but this takes some time to happen.

The initial population fluctuations follow the same kind of pattern in both the continuous space and grid models. The densities of each species over a longer timeframe are graphed in figure 2.3. The densities are more variable in the grid model than they are in the continuous space model for any particular



Figure 2.4: Average species densities at equilibrium for 8 sets of invasion rates, using the continuous space, grid and mean field models. Each colour corresponds to one set of rates. For the continuous space model, the densities are the average over 10,000 density samples taken between 50 million and 100 million interactions in a 150×150 space. For the grid model, densities are the average over 5000 samples taken between 1 and 2 billion interactions on a 500×500 grid.

population size.

But although the densities fluctuations have a higher amplitude, the long run averages of the species densities are similar in both models. Figure 2.4 shows the long-run densities of both models for different invasion rates, along with the fixed points of the mean field equations. Both the grid and the mean field approximations tend to predict larger densities of the least populous species, and consequently they will predict that the risk of ecosystem collapse is smaller than the risk implied by the continuous-space model.

Frean and Abraham (2001) showed that the spatial organisation of species on the grid have distinctive qualities for different sets of invasion rates. There is a much greater degree of clumpiness at equilibrium when the invasion rates are uneven than there is when rates are well-matched.

Three of the most distinctive patterns are shown in figure 2.5. The grid has similar spatial patterns to the continuous-space model in the first two cases, when either all three species are evenly matched, or when one species is much slower than the other two. In the third example, where one species is much



Figure 2.5: Patterns of clustering of species at equilibrium for three sets of invasion rates. Continuous space simulations are shown on the left (150×150 space), and corresponding grid simulations on the right (300×300 grid).

faster than the other two, the landscapes have a similar pattern but different 'scales', because the clusters are much larger in the grid model. This is partially a result of using an exponential distribution, rather than a purely local rule, in choosing the parent-child distance in the continuous space model. A more detailed account of this phenomenon is given in chapter 3 when discussing SCA models with larger, but still partially local neighbourhoods.

2.3 Long-range dispersal

When the SCA model of RPS described in section 1.2 is altered so that interactions are long-range, rather than local, the resulting model is just a pool of individuals with no spatial structure, who interact with one another at random (see section 1.3, page 6). When a simulation is performed using this model, oscillations in the species densities tend to increase in amplitude over time, usually leading to the extinction of the predator of the slowest invader, and subsequently to a monoculture of the slowest invader (Frean and Abraham, 2001).

In continuous space, with a variable population size, the instability seen in the SCA model does not occur.

2.3.1 Continuous-space model with long-range dispersal

The easiest way to introduce long-range dispersal into the continuous-space model described in section 2.1 is to replace the function for choosing the distance between parent and child shown in figure 2.1. The function can be changed so that the child is born at a random location, chosen uniformly over the entire space. However, such a model does not exhibit the instability in species densities which is a feature of the 'long-range' version of the grid model.

It turns out that long-range dispersal is not the primary cause of the instability effect. The instability occurs only if every individual has an equal probability of being selected for an interaction. This is not the same thing as long-range dispersal, for two reasons:

(1) In the continuous space model, a single individual can interact with more than one individual at once. At step 3 in the definition of the model rules (on page 17), a single child individual can compete against

2.3. LONG-RANGE DISPERSAL



Figure 2.6: The frequency with which a child overlaps (left hand side) and invades (right hand side) zero, one, two, three, four or five existing individuals, in the local and long-range versions of the basic continuous-space model described in section 2.1. An overlap with five is the maximum possible. The frequencies were obtained by running a simulation (in a 45×45 space and invasion rates (1,0.6,0.4)) for an initial 500,000 interactions (to ensure that the space is filled up), and then counting the overlaps and invasions during one million subsequent interactions.

multiple existing individuals.

(2) Individuals are not located perfectly evenly over the space. So even if the continuous space model were modified to ensure that interactions could occur only between two individuals at once, some individuals will still be less likely to be chosen, simply because they are in more crowded regions of the space.

The effect of the two factors above is clarified by the consideration of a set of models separating the effects.

2.3.2 1-1 versus 1-n interactions

Firstly, consider a simple change in the specification of the model that forces interactions to be between two individuals at the most. Between steps 2 and 3, the child takes all the individuals it overlaps with, and selects only one of these individuals, at random, to compete against, by 'moving' so that it completely overlaps this competitor.

While this change ensures that interactions are at most 1-1, it also has the following 'side-effects':

(1) The total population can never go down, because existing individuals cannot be removed, only replaced. (The total population can go up,

though, because children can still land in empty space.)

(2) Once the space is so full that there are no empty spaces, the number of individuals in any subregion of the space can never change, because there is no movement of individuals once they are initially placed.

The second side-effect means that even in the long-range dispersal version of this model (where children land with equal probability on all parts of the space), existing individuals in emptier regions, with larger 'basins of attraction', will be interacted with more often than those in close-packed regions of the space.

Enforcing interactions to be 1-1 has a significant effect on the system. Figure 2.6 shows, for the original continuous-space model, how often children land on zero, one, or multiple existing individuals (once the space fills up), and how often they succeed in invading those individuals. 1-2 interactions are the most common type in both the local and long-range dispersal models, although successful invasions are much more likely to come from 1-1 interactions because of the difficulty of successful competition against two individuals (probability r^2 for invasion rate r), and also because of the small probability that both individuals happen to be of the child's prey species.

Successful invasions against two or more individuals are more likely in the local model because the localised invasion entails that nearby individuals are quite likely to be the same species (if they are different species, the rules dictate that no invasion can occur). On the other hand, successful invasions against a single individual are less likely in the local model because nearby individuals are likely to be from the same species as the invader.

In both models, there are a significant number of successful invasions against zero individuals (children falling on empty space). The continual presence of empty space in the system is what makes the 1-n models so much more stable than the 1-1 models (see the top two rows in figures 2.7 and 2.8). This is because empty space is always invadable by all three species regardless of their invasion rates or current densities, so it provides a 'refuge' into which weak or threatened individuals will reproduce.

The stabilising effect of a fluctuating population is not unique to RPS ecosystems. A similar effect has been previously noted by (Mitteldorf and Wilson, 2000), in which a fluctuating population with empty space prevents the extinction of altruists in competition against selfish individuals.


Figure 2.7: Early density fluctuations in the 1-*n*, 1-1, and fixed population models for very uneven invasion rates (1.0,0.2,0.3). The 1-*n* models (top row) are the basic continuous-space model of section 2.1, and the long-range variant described in section 2.3.1. The 1-1 models in the second row are the local and long-range versions of the 1-1 model described in section 2.3.2. All these simulations start with populations of 6000 in a space of 100 × 100. The 'local, fixed neighbours', and 'long-range, uniform in pop' models are equivalent to the SCA model from section 1.2 (81 × 81 grid) and the 'pool of individuals' model from section 1.3 in which every individual has an equal probability of being selected (fixed population of $81 \times 81 = 6561$).



Figure 2.8: Long-run fluctuations for the six models described in figure 2.7 with invasion rates of (1.0,0.4,0.6), and initial populations and space as in figure 2.7.

2.3.3 Fixed populations and equal interaction probabilities

As I have mentioned above, simply enforcing 1-1 interactions using the above method does not reproduce the extinctions observed by Frean and Abraham (2001) in fixed populations with long-range dispersal. The crucial difference between the two models is that in the Frean and Abraham version, individuals are chosen uniformly from the entire population, and this is not true of the 1-1 model described above. Although the 1-1 model chooses the child's loca-

2.4. DISCUSSION

tion uniformly, and then chooses the victim uniformly from the set of overlapping individuals, this does not amount to choosing the victim uniformly from the population because some regions are more closely packed than others.

While this may seem like a small difference, it has a profound effect on the stability of the system that can be seen in figures 2.7 and 2.8. The middle right diagrams show the behaviour when victims are taken from the set of individuals overlapping the child, and the bottom right diagrams the behaviour when victims are chosen uniformly from the entire population. Although there is little difference between the two in the early stages (figure 2.7, in the long run choosing from all individuals with equal probability leads to increased fluctuations and eventually an extinction (figure 2.8).

This is because equal interaction probabilities lead to density oscillations that are truly global in space and time. Oscillations spread out to all regions simultaneously, which means that all areas are 'in phase'. Unequal interaction probabilities, on the other hand, effectively produce a small time delay between regions, because the more probable individuals tend to be invaded a little bit earlier than the less probable ones. The delay is just long enough to keep the oscillations out of phase so that the dominant species in the low probability regions is different from the dominant species in the high probability regions at any particular instant.

2.4 Discussion

When dispersal is localised, the SCA and continuous space models of RPS have the same general characteristics. The grid model predicts slightly bigger density oscillations than the continuous space model, but it also predicts slightly closer average densities for the three species. These two effects will cancel somewhat when the risk of an extinction is being evaluated.

There is a big difference between the two models when dispersal is longrange. This is because in continuous space, dispersal can be long-range and the probabilities of interacting can vary from individual to individual because they may have different sized 'catchments'. This is not possible with longrange dispersal on a grid-based model, because all individuals' catchment areas must be the same size on any lattice.

It turns out that non-uniformity in interaction probabilities tends to stabilise the system, and even the small amount of nonuniformity in quite closelyspaced individuals (for example in the bottom right of figure 2.2) causes sufficient stability to eliminate the likelihood of extinctions.

Whether or not extinctions occur in real ecosystems will depend on the relative rates of movement and interaction. The model here has presumed that movement and interaction occur at exactly the same rate. The laboratory experiments of Kerr et al. (2002) observed extinctions in a RPS system with longrange dispersal, but the rate of movement was very high. They mixed three strains of *E. coli* (in a cyclic dominance relationship) in a flask that was shaken at a rate of 125 revolutions per minute. This rate of movement may have been sufficient to ensure that interaction probabilities were effectively uniform.

Real ecosystems in earthly space are very likely to contain at least the small amount of nonrandom interaction present in the long-range, continuous space model. The implication for real RPS ecosystems is that we should expect them to be very stable, even when individuals are able to mix relatively freely.

Chapter 3

The effect of grid size and intermediate-range dispersal on extinction in spatial RPS models

RPS ecosystems with local interactions exhibit a high degree of long-term stability in community composition, while those with random interactions tend to be unstable until two species become extinct. This effect has been shown both for real bacterial ecosystems (Kerr et al., 2002), and in cellular automata simulations (Frean and Abraham, 2001).

The two situations which give these divergent results can be generalised into a single system with one parameter. This is because random interactions and local interactions are two end-points of a continuum of dispersal distance. At one end, a stochastic CA with random interactions between all individuals in the population models an ecosystem in which offspring are equally likely to disperse any distance or in any direction across the space. An example would be an ecosystem made up of plants whose seeds may be blown over very long distances. At the other end, a CA with local interactions is better suited for the modelling of plants which only grow by sending out runners into adjoining space. In between are organisms such as most animals, which can move longer distances with some maximum range.

There is a big difference between the outcomes at the two endpoints, and the subject of this chapter is the nature of the transition between them. Specifically, to answer the question of whether there is a gradual change towards instability as the dispersal distance increases, or whether there is some sort of critical threshold.

We also know that larger RPS grids tend to be more stable than smaller ones, so a secondary purpose is to find out whether a similar sort of transition from stability to instability arises when decreasing the size of the grid rather than increasing the dispersal distance.

In this chapter I describe the results of CA simulations that are intended to answer these questions. The findings of these experiments are:

- (1) RPS systems become increasingly stable as the square grid is increased in size, and the stability increases at a faster than exponential rate with respect to the grid size (when measured by the length of one side of the square).
- (2) RPS systems on a grid with a particular size become more stable as the radius of the local neighbourhood is decreased, and the stability increases at a faster than exponential rate with respect to decreasing neighbourhood radius.
- (3) The grid size and neighbourhood size variables affect stability independently of one another, and a linear combination of the two can predict the stability of the system.
- (4) Although very long-range dispersal leads to extinctions, dispersal over short distances can increase stability in comparison to purely local dispersal.

The rest of the chapter proceeds as follows. First I describe some details of the stability measure used in the experiments. The effect of varying the size of the grid is described in section 3.2, the effect of varying the interaction neighbourhoods is described in section 3.3, and section 3.4 shows the combined effect of variations in both variables. The results are discussed in section 3.5.

3.1 Time to extinction (TTE) experiments

For the experiments described in this chapter, I have used a cellular automaton model of the RPS system, with rules identical to the one described in section 1.2 on page 4. To assess the stability of the model ecosystem with a particular set of parameters, I just count the number of generations until an extinction occurs. This number, which I call the time to extinction or TTE, quantifies an ecological concept similar to 'persistence stability' as defined in

3.2. INCREASING THE GRID SIZE

Johnson and Mann (1988) and Johnson and Boerlijst (2002), a concept which is explained in more detail in chapter 10. Persistence stability is usually not concerned with the magnitude of changes in species densities over time; it just measures whether species population sizes remain non-zero (Connell and Sousa, 1983).

In a CA, the time during which all three species survive is a good measure of the stability of the ecosystem, and it is easy to calculate for a particular simulation, but it has the drawback of making it difficult to get a result for very stable ecosystems.

Coexistence, extinction and stability in CA simulations

It will never be possible to say that coexistence in a particular CA simulation is completely stable, because in a stochastic system it always remains possible for some freak chain of events to bring about an extinction. All sites in a CA grid are connected to at least four others, so it is theoretically always possible to reach the state in which the population-*N* grid is saturated by species *A* after only $(1 - \rho_A)N$ interactions. However, as the *N* increases this probability of such an extinction becomes vanishingly small.

It is feasible to run simulations for a large number of generations and see whether extinctions occur, and that is the approach I have taken in these experiments. This allows systems of various sizes to be compared for stability, but it is not possible to find a grid size at which a given system becomes indefinitely stable using this method.

3.2 Increasing the grid size in a local dispersal system

In the experiments detailed here spatial RPS simulations are run for many different sized grids to find the nature of the transition from instability to stability.

3.2.1 Two reasons why large grids increase stability

There are two reasons why larger grids should foster coexistence. The first reason is not related to RPS ecosystems. It's simply that even unfit species can

be expected to survive longer in larger populations given the same species densities. The unfit species starts out in greater numbers in the large population and survives longer through luck alone. For example, the absolute minimum time to extinction of $(1 - \rho_A)N$ described in 3.1 above scales with the population size *N*.

The second reason is more interesting and relates to the apparent characteristic cluster sizes seen in spatial RPS simulations with various invasion probabilities, examples of which are shown in figure 2.5 (page 21). An intuitive argument states that if there is a characteristic maximum cluster size for a particular set of species, then grid sizes at or below this maximum cluster size are likely to result in extinctions, while those well above the maximum cluster size may be able to remain in a state of coexistence.

3.2.2 RPS versus random drift

These two effects are not easy to distinguish if time to extinction is used as the only measure of stability. In order to separate the two effects, I will compare the RPS system to a second system, which I call the 'drift' system, in which individuals are able to invade one another at random, regardless of species.

In populations consisting of multiple types of individual each with equal fitness, it has been shown that all but one type will eventually go extinct through the process of genetic drift (see for example the explanation given by Suzuki et al. (1989)). Dick and Whigham (2005) have modelled the genetic drift process in a spatially structured environment and found that in the absence of selection, even extremely spatially structured populations will evolve by genetic drift to a state of extinction.

I will compare the RPS system to an interacting population with neutral fitness effects in order to show the extent to which selection under the RPS rules is responsible for the observed stability. The drift system should show coexistence caused by the first of the two effects described in section 3.2.1 above, but not the second effect, because there should be no maximum cluster size. The RPS system, on the other hand, will exhibit both effects, so by subtracting the coexistence times found in the drift system, the stability effect due to RPS can be found.

The CA simulations I provide below are asexual models for the purposes of simplicity and so the system I am calling 'drift' is different to the genetic drift in sexual populations studied by Dick and Whigham (2005).

3.2.3 Invasion rates and averaging of TTE

In the simulations in this and the next section I give each species an invasion rate of 1 against all other species in the drift system, and against the next species in the cycle for the RPS system. Invasions involving two members of the same species have no effect, and can be ignored, but because the sum of all the invasion rates (for interactions involving different species) is twice as high in drift as it is in RPS, there will be twice as many invasions on average in the drift system than there are in the RPS system.

Because of the large variation between simulations with identical initial states, TTEs over multiple simulations must be averaged out in order to get useful comparisons.

3.2.4 TTE increases with grid size

Figure 3.1 shows TTEs for thousands of simulations with gradually increasing grid sizes for the RPS and drift systems. With each increase in the grid size, more and more of the CA simulations run for longer before an extinction occurs. The average TTE, shown by the black lines, increases sharply in the RPS system (fig. 3.1(a)) and only slowly in the drift system (fig. 3.1(b)).

3.2.5 Trends in average TTE

The character of these increases in average TTE is shown more clearly in figure 3.2, which also shows that the average TTE is far longer in the RPS system than it is in the drift system for all the grid sizes (see the scale on fig. 3.2(b)).* Stability in RPS continues to climb steeply while the stability of the drift line appears to be levelling out, indicating that the stability rise with grid size is faster than exponential for RPS, but merely linear for the drift system.

Finally, as the grid size increases, the average TTE in the drift system becomes very insignificant when compared to the extinction time in the RPS system, indicating that increased coexistence caused purely by a larger population becomes insignificant compared to coexistence caused by the characteristic maximum cluster size of stable RPS systems.

^{*}As discussed above the drift system as implemented here has an effective dynamical timescale that is twice as fast as that of the RPS system. But even when these different timescales are accounted for, RPS still has a longer TTE for all the grid sizes displayed in the figure.



Figure 3.1: Time to extinction for 10,000 simulations at each of ten different grid sizes: 21×21 , 22×22 , ..., 30×30 . (a) is a RPS system with equal invasion rates and (b) is the three-species drift system in which all species have an equal probability of invading one another. Each blue dot marks the TTE for a single simulation; the red dots and black line show the mean TTE for the 10,000 simulations at each grid size. The mean is not shown for the RPS system with a 30×30 grid because 28 of the 10,000 simulations had not reached an extinction after the maximum 2,000,000 generations.

3.2.6 Distribution of TTE

Figure 3.3 shows the distribution of TTEs of the simulations performed above. The distributions all exhibit a characteristic long tail.

The drift distributions show a clear peak which appears to be lacking in the case of RPS, but this doesn't mean that only the drift system has a characteristic extinction time: in fact a small peak is visible even for RPS upon examination of the extinctions which occur early in the simulations, examples of



Figure 3.2: Average TTE for the RPS (red, right hand axis) and random drift (blue, left hand axis) systems, shown on (a) linear and (b) log scales.

which are shown in figure 3.4. For these examples, RPS extinctions are less likely to occur in the first 150 generations than they are at some later times, so even the RPS distributions are peaked.

The shape of the distributions shown in figure 3.3 do not appear to change with respect to the grid size, even though they are stretching out over longer times. (Although the distribution of extinction times in RPS with a grid size of 30×30 looks different to those from the smaller grid sizes, this is probably because 28 simulations were cut from the distribution because their values exceeded the maximum simulation length of 2,000,000 generations).

Figure 3.5 shows that in the first 150 generations, extinctions are actually slightly more likely in RPS than in the drift system. However, as figures 3.3 and 3.4 showed, if one of these extremely early extinctions does not occur, then extinctions are less likely under RPS for ever after.

3.3 Decreasing the dispersal range

The effect of the dispersal distance, for a given grid size, can be checked by running CA simulations in which the interaction neighbourhood is altered. As discussed in the introduction to this chapter, in all the CA models presented so far it has been assumed either that an individual can only interact with its immediate (north, south, east, west) neighbours on the grid, or that an individual can interact with any other individual in the entire population. These two scenarios are special cases in which the neighbourhood radius is



Figure 3.3: Distribution of TTE for the RPS and 'drift' systems, where each bar shows the number of simulations (out of 10,000) whose TTEs fall into one of 70 equally spaced bins. The overall shape of the distributions remains relatively constant with respect to the system size.



Figure 3.4: Close-up of distributions of TTEs occurring in the first 2500 generations. For each of the four cases, 10,000 simulations were run, and the numbers in brackets indicate how many of those 10,000 simulations had extinctions after the 2500th generation.



Figure 3.5: Number of extinctions out of 10,000 simulations which occur within the first 150 generations.

one or infinite. In between these extremes I have used increasing radii, three examples of which are shown in figure 3.6.

A large neighbourhood radius reduces clustering and tends to homogenise the grid state, as can be seen in the grid snapshots in figure 3.7.

3.3.1 The effect of neighbourhood size on TTE

In the case of dispersal distance it is the small neighbourhood radii which lead to the most long-term stable coexistence. Therefore in these simulations I show the effect of decreasing the neighbourhood size from some large number



Figure 3.6: Examples of three neighbourhood sizes, or dispersal distances. The individual in the black site can interact with all the sites shown in grey.



Figure 3.7: The state of a 100×100 grid after 4000 generations for three sets of invasion rates and three neighbourhood radii. Rocks are shown in black, scissors in grey, and paper in white, and the invasion rates are listed in order rock, scissors, paper.



Figure 3.8: Effect of decreasing neighbourhood size on average TTE (over 100 simulations) for the RPS and drift systems on 100×100 and 80×80 grids. For RPS, stability increases dramatically when the neighbourhood size falls. For the drift system, neighbourhood size does not affect stability.

for which the average TTE is similar to the average TTE for infinite dispersal.

Figure 3.8 shows average TTE for the RPS and drift systems when the neighbourhood radius is reduced from 25 down to 13, on two different grid sizes. The leftmost column on the *x*-axis, marked 'Inf', shows the TTEs for an infinite dispersal range in which the entire grid is covered. In the RPS system TTE increases at a faster than exponential rate with respect to decreasing neighbourhood radius.

For the drift system, the neighbourhood size appears to have no effect on coexistence, providing evidence for the hypothesis that the duration of coexistence in the drift system depends solely on the size of the population. This is another good reason to believe that the drift system results are an appropriate way to separate large-population caused coexistence from any inherent stability encoded in the RPS invasion rates.



Figure 3.9: Distributions of TTE for 10,000 RPS simulations on 100×100 grids, for three neighbourhood sizes: (a) full distributions showing all samples; (b) frequency of simulations with TTE less than 2000 generations.

3.3.2 Neighbourhood size and TTE distributions

A decrease in neighbourhood size affects the shape of the distribution of TTE, an outcome that is not observed when changing only the grid size. The full distributions, shown in figure 3.9(a) stretch out in response to decreasing neighbourhood radius.

Distributions over the first 2000 timesteps (fig. 3.9(b)) indicate the presence of a minimum time before which extinctions are extremely unlikely. Unlike the mean TTE, this minimum barely shifts in response to decreasing neighbourhood size. Even with infinite dispersal, when we start with equal densities, the three species populations will go through several cycles before any of them drop to zero.



Figure 3.10: Variation of TTEs (average over 20 simulations, up to a maximum of 5,000,000 generations) with grid size and neighbourhood size, when $r_r = r_s = r_p$. The shading is linear with extinction time with pure black representing extinction by 100 generations through to pure white representing the maximum 5 million generations.

3.4 Combinations of grid size and dispersal distance

Although there can be no hard stability threshold, for reasons discussed earlier, the dramatic rise in stability indicates an effective threshold for both grid size and dispersal distance which is shown in figure 3.10. The diagonal separating the black region of early extinction and the white region with stable coexistence depends on both grid size and dispersal distance in a near-linear combination.

3.4.1 Relative importance of grid size and neighbourhood

The slope of the boundary line in figure 3.10 suggests that a decrease of one in neighbourhood radius is roughly equivalent, in terms of stability, to an increase of five in the grid size.



Figure 3.11: Variation of TTEs (average over 20 simulations, up to a maximum of 5,000,000 generations) with grid size and neighbourhood size, for $r_r = 1$, $r_s = 0.5$, $r_p = 0.5$ (top), and $r_r = 1$, $r_s = 1$, $r_p = 0.3$ (bottom).

3.4.2 Effect of invasion rates and cluster size

I suggested that the characteristic maximum cluster size might be a determining factor in whether or not a particular grid size or neighbourhood size could bring about stability or extinction in the spatial RPS game. Figure 3.11, which shows a similar threshold effect to the one in figure 3.10 for different sets of invasion rates, provides evidence in favour of this claim.

Figures 3.10 and 3.11 show the effect of the grid size and dispersal parameters for the systems displayed in figure 3.7, and it turns out that stability does in fact decrease with increases in the cluster sizes. A crude measure of clustering can be obtained by taking a few sites at random from a grid and for each of these sites, counting the number of sites in one direction that are in the same state. Table 3.1 shows the results of doing this for simulations using the three sets of invasion rates from figure 3.7.

The third column of the table shows that the general trend towards instability is quite well predicted by the amount of clustering in the grid with local dispersal.

3.4.3 Stability is maximised for small, but not minimum neighbourhoods

Cluster sizes do not fully predict the stability of the system, because for all three sets of invasion rates in figures 3.10 and 3.11, systems with a very small amount of non-contiguous dispersal (neighbourhood radius 2-3) are in fact more stable than those with purely local dispersal (neighbourhood radius 1). A comparison of the first and second rows of figure 3.7 shows that radius 3 systems are typically less clustered than radius 1 systems.

As neighbourhoods get even larger, clustering increases even further (see the bottom row of figure 3.7), but now stability decreases rather than increases. There is therefore no simple correspondence between cluster size and stability.

The extra stability brought about by a small amount of non-local dispersal is caused by the ability of a species to leap over regions occupied by its 'predator' into a region occupied by its 'prey', which is impossible for pure local dispersal. With pure local dispersal, large areas of species *B* that become surrounded by their predator species *A* will normally face local extinction (except in the case where their prey species *C*, pursuing the predator, breaks through into areas occupied by the original species *B*). The extra mixing provided by non-contiguous dispersal reduces the number of these local extinctions by

Inva	asion	rates	Cluster measure	Smallest stable grid				
1.0	1.0	1.0	3.2	36×36				
1.0	0.5	0.5	4.0	39×39				
1.0	1.0	0.3	5.5	54 imes 54				

Table 3.1: Clustering for RPS simulations with three sets of invasion rates. For each set of invasion rates, the measurement given in the second column is calculated in the following way. First, 100 simulations were run for 4000 generations on a 100 \times 100 grid with purely local dispersal. For each of these grids, 100 distinct sites were chosen at random from all over the grid, and for each random site, the number of contiguous sites to its right occupied by the same species as the original site were counted. Each number in the table is therefore the average over 10,000 sites from 100 different grids. The third column shows the smallest grid size from the simulations in figures 3.10 and 3.11 for which all of the simulations reached the maximum 5 million generations.

creating more alternative paths through the space.

The effect is more pronounced for uneven, more clustered environments, shown by the especially large jump in stability between radius 1 and radius 2 neighbourhoods for the most clustered system shown at the bottom of figure 3.11. This is because the most clustered, uneven systems have small densities of very fast invaders which come through in thin waves that are easy to jump over with the addition of a very few non-local interactions.

3.5 Summary and discussion

In the spatial RPS game, there is a simple correspondence between the size of the grid and the length of time before extinctions occurs. For a set of invasion rates, the stability of the system grows faster than exponentially with every increase in the size of the grid, and at any given timescale there will be an effective threshold above which the system can be considered stable because extinctions become vanishingly unlikely.

There is a similar, independent stabilising effect caused by neighbourhood size in spatial RPS systems with non-local dispersal. Large dispersal distances tend to amplify fluctuations in species densities that lead to global extinctions, but these fluctuations decrease once the neighbourhood size becomes sufficiently small. Subsequent decreases in neighbourhood size tend to cause a greater than exponential increase in stability.

3.5. SUMMARY AND DISCUSSION

However, for a critical range of grid sizes, the system can become less stable again when the neighbourhood size becomes very small, because of an increase in extinctions at the local level.

This last result is especially surprising because it has generally been believed that spatial structure provides a mechanism for the maintenance of diversity (Hassell and May, 1973; Chesson, 2000; Durrett and Levin, 1994a; Hassell et al., 1994; Tilman and Kareiva, 1997; Dieckmann et al., 2000). The effect has been noted in RPS (Johnson, 1997; Durrett and Levin, 1998) and in many other systems such as the spatial Prisoner's Dilemma (Nowak and May, 1992), and more general models of evolutionary altruism (Mitteldorf and Wilson, 2000). The result presented here shows that the effect of spatial structure is more complicated, because there are cases in which the addition of spatial structure causes a decrease in system diversity.

Part II

Approximations of spatial structure

Chapter 4

Pair and local structure approximations for rock-paper-scissors ecosystems

The equilibrium behaviour of rock-paper-scissors ecosystems depends crucially on whether the population is spatially structured or well-mixed. Because it is rare for natural populations to be completely well-mixed, it is important to have good models for the more common situation where interactions are non-random. Cellular automata help to fill this role, but they are computationally expensive in comparison to mean field models using ordinary differential equations. The amount of computation required can make the task of exhaustively exploring the space of parameters and initial conditions into quite a painstaking process.

Local structure approximations are a method for including a small amount of spatial information in an ODE model that is computationally efficient to simulate numerically, and in this chapter I describe the application of the technique to RPS ecosystems. Because these models provide more spatial information than the mean field, we should expect them to produce more accurate models. However, I show that this is not always the case for RPS systems, and that some of their properties are better predicted using the the simpler mean field model.

4.1 Pair and local structure approximations

Pair approximations originated in statistical mechanics, and were first used for ecological models by Matsuda et al. (1987). They are a class of models which include a limited amount of spatial information in a dynamical system by approximating a limited number of the local interactions that would take place in a full cellular automaton model.

The general technique is to track correlations between the types of neighbouring pairs of individuals. So while the mean field model described in section 1.1 consists of equations for changes in species densities, a pair approximation model of the same system specifies equations that track the densities of neighbouring pairs of individuals. Good examples of ecological pair approximation models are given in Harada et al. (1995) and Matsuda et al. (1992). 'Local structure approximations' is a name given to the more general class of models in which neighbouring groups larger than just pairs may be involved (Hiebeler, 1997).

Although pair and local structure approximation models assume continuous population densities, like the mean field models, they also assume some level of discreteness of space, like the SCA models, because neighbour relations must be represented explicitly. They are always significantly more complex than corresponding mean field models, but very much simpler than stochastic cellular automata, and in many cases have been successful in describing the essential effects of spatial structure without the overhead of agent-based models

The use of the pair approximation technique has been shown to improve significantly on the mean-field model for some populations (Sato and Iwasa, 2000; van Baalen, 2000; Hiebeler, 2000; van Baalen and Rand, 1998; Caraco et al., 2001), although there are also cases in which it fails to improve predictions (Ives et al., 1998; de Aguiar et al., 2004). RPS systems appear to fall into the latter category. In the next few sections I show that for the most important statistical property of a RPS ecology, species composition, two simple local approximations that are relatively easy to derive do not deliver an improvement in predictive accuracy over the spatially deficient mean field model. Szabó et al. (2004) have obtained similar results.

4.2 Application to RPS

RPS ecosystems are especially sensitive to assumptions about discreteness of individuals and spatial structure. As discussed in Frean and Abraham (2001), and shown in figure 1.2 on page 5, spatially explicit models make very different predictions to the mean field model about the stability of the system.

The basis of pair approximations is to think of the population as being on a graph in which an edge between two sites *x* and *y* indicates that individuals occupying those sites are able to interact with one another. For simplicity, and for ease of comparison with the cellular automaton model, I will use a regular grid rather than a general graph.

Because the pair approximation is an approximation to a cellular automaton, it can only be expected to provide more accurate predictions of an ecosystem insofar as the SCA's predictions are accurate. Real ecosystems are usually not lattice-based, and will probably have properties more like those of the models described in Part I of this thesis.

4.2.1 Pair transformation dynamics

In the RPS system the mean field dynamics are shown in figure 4.1(a), where rock sites can turn into paper sites, paper sites can turn into scissors sites, etc., through invasions. The pair dynamics are shown in figure 4.1(b) and (c). Figure 4.1(b) shows the ways in which pairs of sites can change into other pairs if invasions are allowed in both sites at once, and figure 4.1(c) shows the ways in which pairs through the effect of a single invasion in one site.

4.2.2 Symmetry in pair types

There are $3^2 = 9$ possible configurations of three species in two sites, but these are reduced to six different nodes in the diagrams because of the symmetry of the invasion rules: the density of pairs of the form [*ab*] will be the same as that of [*ba*] pairs, so a single node is used to represent both. For the remainder of this chapter I will write the nine pair states where ordering is important in square brackets: {[rr], [rs], [rp], [sr], [ss], [sp], [pr], [ps], [pp]}; and the six unordered states without square brackets: {rr, ss, pp, rs, sp, pr}.

The density ρ_{xy} will refer to the proportion of pairs of which one is in state *x*



Figure 4.1: The rock-paper-scissors system: (a) Dynamics of single sites (b) Dynamics of pairs when invasions at both sites are allowed (c) Dynamics of pairs of sites with invasions at a single site only.

and the other is in state *y*, regardless of the order, and $\rho[xy]$ is the density of pairs where the first is in state *x* and the second is in state *y*. So $\rho_{xx} = \rho[xx]$, and when $x \neq y$, then $\rho_{xy} = \rho[xy] + \rho[yx]$.

4.2.3 Counting of pairs

There are two common approaches to counting the pairs. Rand (1999), for example, presents an account which applies to arbitrary graphs, in which all edges are counted twice, so for two sites x and y, the edge from x to y is distinct from the edge from y to x. This means that for the regular lattice with Q neighbours per site, there are Q times as many pairs as there are sites. When all edges are counted twice, this will ensure that $\rho[xy] = \rho[yx] = \frac{1}{2}\rho_{xy}$, a property that should be expected to hold in the long run in the SCA because the invasion rules are symmetric. Two examples of the ways in which edges are counted are shown in table 4.1.

If the task is simply finding correlations between the states of neighbouring



	ordered pairs,						unordered,				unordered				
	counted in both directions						one direction			densities					
	[rr]	[rs]	[sr]	[sp]	[ps]	[rp]	[pr]	rr	rs	pr	sp	$ ho_{rr}$	$ ho_{\sf rs}$	$ ho_{pr}$	$ ho_{ m sp}$
(1)	2	1	1	1	1	1	1	1	1	1	1	1/4	1/4	1/4	1/4
(2)	4	2	2					2	2			1/2	1/2		

Table 4.1: Two approaches to counting pairs in the 2×2 grids (1) and (2), assuming non-periodic boundaries. When the order of pairs is important and the edges of the graph are counted in both directions, there are 8 pairs; when order is unimportant and edges are counted once, there are only four. The densities of unordered pairs remains the same no matter which method is used.

sites, it is unnecessary to distinguish, for example, between pairs in state [rs] and those in state [sr]. Therefore the unordered pair state descriptions will generally be used in the following sections. When counting unordered pairs, edges can be counted once only, and there will only be Q/2 times as many pairs as there are sites.

4.2.4 Homogeneous and heterogeneous pairs

It can also be seen from the diagrams in figure 4.1 that structurally there are really only two different types of pairs, the '*aa*-type' pairs, in which each site of the pair is occupied by the same species (rr, ss, and pp), and the '*ab*-type' pairs in which the two cells are occupied by different species (rs, sp, and pr). In figure 4.1, all three homogeneous or *aa*-type pairs occupy similar positions in relation to the other pairs, as do the three heterogeneous or *ab*-type pairs.

The development of the model in the following section will make use of the similarity within the two pair types, because it can be used to reduce the number of equations.

4.3 Derivation of pair correlation equations

Figure 4.1(c) shows the possible state changes if invasions are allowed to occur at only one of the two sites in a pair of neighbours. Using only single invasions at one site, we can model a system similar to the SCA simulation where at each timestep, we pick two neighbouring sites at random from a grid, see if the first successfully invades the second, and if so, update any pair densities which are affected. The pair densities that will change after a successful invasion are the densities of those pairs in which at least one of the pair is in the same state as the invaded site.

4.3.1 Rate of change of homogeneous pairs

From figure 4.1(c), we can see that the *aa*-type pairs all have only one arrow leading into them and one arrow leading out, corresponding to the ways in which these pairs can be created and destroyed from other pairs. The corresponding equation for the rate of change of each of the *aa*-type pairs will have two terms, each corresponding to an arrow in the diagram.

For example, with the rate of change of rr pairs, there will be two terms in the equation: a positive term corresponding to the arrow going out of rr to pr, and a negative term corresponding to the arrow going into rr from rs. The resulting equation is of the form

$$\frac{\mathrm{d}\rho_{\rm rr}}{\mathrm{d}t} = R[\mathrm{rs} \rightarrow \mathrm{rr}] - R[\mathrm{rr} \rightarrow \mathrm{pr}]$$

where $R[rs \rightarrow rr]$ is the rate at which rr pairs are created when a rock invades the scissors in a rs pair, and $R[rr \rightarrow pr]$ is the rate at which rr pairs are destroyed (and turn into pr pairs) when one rock in an rr pair is invaded by a paper.

In general, for all of the homogeneous pairs, there is an equation for the rate of change. So we have three equations represented by

$$\frac{\mathrm{d}\rho_{aa}}{\mathrm{d}t} = R[ab \to aa] - R[aa \to ca] \tag{4.1}$$

for all $a, b, c \in \{r, s, p\}$ such that b = prey(a) and c = predator(a), where

$$prey(r) = s$$
, $prey(s) = p$, $prey(p) = r$, and $predator(r) = p$, $predator(s) = r$, $predator(p) = s$.

Equation (4.1) is just a general way to represent the rate of change of all three *aa*-type pairs using *a*, *b*, and *c* to represent the three species.

4.3.2 Rate of change of heterogeneous pairs

Similarly, the rate of change of the heterogeneous *ab*-type pairs has four terms corresponding to the four edges which terminate on the *ab*-type nodes in figure 4.1(c):

$$\frac{\mathrm{d}\rho_{ab}}{\mathrm{d}t} = R[bb \to ab] + R[ca \to ab] - R[ab \to bc] - R[ab \to aa] \tag{4.2}$$

for all the same values of *a*, *b*, *c* as in equation (4.1).

4.3.3 Rates of transformation of pairs

The terms of the form $R[wx \rightarrow yz]$ in equations (4.1) and (4.2), which describe the rates of transformation of pairs in state wx to pairs in state yz, need to be expressed using the invasion rates r_x of single sites, and pair densities ρ_{xy} .

There are only three kinds of transformation when only one invasion is allowed: $R[ab \rightarrow aa]$, $R[aa \rightarrow ca]$, and $R[ca \rightarrow ab]$, when b = prey(a) and c = predator(a). I will discuss an example of each type.

Case 1: $ab \rightarrow aa$

An example of $R[ab \rightarrow aa]$ is $R[rs \rightarrow rr]$. The only way in which rr pairs can be created is when an r invades the s of an rs pair; any invasion by an s or a p cannot increase the density of rr pairs.

The overall rate will be determined by the proportion ρ_{rs} of pairs in state rs and the rate r_r at which the r successfully invades s in these pairs. If the procedure were to occur in the same way as the simulations described in section 1.2, we could think of the ρ_{rs} as the chance that the random potential invader and neighbouring victim picked are r and s. The required selection and invasion that will produce rr will only happen $\rho_{rs}r_r$ of the time.

When the invasion is successful, one rr pair is created from the initial rs pair, but there may be additional rrs created, because the s that was invaded has Q - 1 other neighbours, and each of those neighbours might also have been in state r. The notation $q_{x|yz}$ is used to represent the probability that any one of the Q - 1 neighbours of the y in a yz pair is in state x. The expected number of rr pairs created by the invasion of an s by an r is therefore $1 + (Q - 1)q_{r|sr}$, and the overall transformation rate is

$$R[\mathsf{rs} \to \mathsf{rr}] = \rho_{\mathsf{rs}} r_{\mathsf{r}} (1 + (Q - 1)q_{\mathsf{r}|\mathsf{sr}}). \tag{4.3}$$

Case 2: $aa \rightarrow ca$

 $R[aa \rightarrow ca]$ is simpler. For example, consider $R[rr \rightarrow pr]$. The invading p has to come from somewhere, which means that one of the r's must be part of a pr pair before the invasion, so we only need to consider a successful invasion by a p when a pr pair was initially selected, and this happens $\rho_{pr}r_p$ of the time. The expected number of prs created in this case is $(Q - 1)q_{r|rp}$, so the overall rate is

$$R[\mathsf{rr} \to \mathsf{pr}] = \rho_{\mathsf{pr}} r_{\mathsf{p}} (Q-1) q_{\mathsf{r}|\mathsf{rp}}. \tag{4.4}$$

Case 3: $ca \rightarrow ab$

Finally, consider $R[pr \rightarrow rs]$, an example of $R[ca \rightarrow ab]$. Here the required transformation can only occur when an s invades a p, which will occur $\rho_{sp}r_s$ of the time, and will create an expected $(Q-1)q_{r|ps}$ new rs pairs, so

$$R[\mathsf{pr} \to \mathsf{rs}] = \rho_{\mathsf{sp}} r_{\mathsf{s}} (Q-1) q_{\mathsf{r}|\mathsf{ps}}. \tag{4.5}$$

4.3.4 The pair closure

The transformation rates $R[wx \rightarrow yz]$, on a true grid, are dependent on all the neighbouring cells. If for example cells interact with their four closest neighbours (Q = 4), then the relevant cells for all interactions with the original pair [wx] are all of the cells

	y	Z	
y	w	x	Z
	y	Z	

A *y* could invade *w* from any of the three cells marked *y* above, and a *z* could invade *x* from any of the cells marked *z*, so the transformation rate $R[wx \rightarrow yz]$ depends on the densities of all the ordered triples $\begin{bmatrix} y \\ wx \end{bmatrix}$, [ywx], $\begin{bmatrix} wx \\ y \end{bmatrix}$, and the triples $\begin{bmatrix} z \\ wx \end{bmatrix}$, [wxz], $\begin{bmatrix} wxz \\ z \end{bmatrix}$.

Terms like $q_{y|wx}$ and $q_{z|xw}$ are used to indicate that we don't care which of the w's neighbours is in state y, or which of the x's neighbours is in state z. All of these $q_{x|yz}$ terms can be rewritten using only pair densities if we assume the 'pair approximation', which is referred to by Rand (1999) as the 'pair closure with Bernoulli trials'. To avoid having to describe these conditional probabilities in terms of the densities of triples, the $q_{x|yz}$ terms are approximated by

$$q_{x|yz} \approx q_{x|y}$$

56

which claims that the chance of one of the neighbours of the y in an yz pair being in state x is approximately the same as the chance that a neighbour of a y is in state x.

It is easy to see why this assumption will often be unrealistic. If *xs* disperse locally throughout the grid, for example, then we would expect to see some local clustering of *xs*, and $q_{x|yx}$ would probably be significantly higher than $q_{x|y}$. In general, the greater the clumpiness, the less realistic the pair approximation will be. But the pair approximation is necessary to close off the set of equations, otherwise the rate of change of triples will need to be calculated, which would themselves depend on the densities of quadruples and even larger clusters of sites.

The pair approximation allows the densities of triplets to be rewritten as

$$\rho[xyz] \approx \frac{\rho[xy]\rho[yz]}{\rho[y]} = \frac{\rho[xy]\rho[yz]}{\sum_i \rho[yi]}$$

The conditional probabilities in the pair transformation rate equations can be rewritten as

$$q_{x|yz} \approx q_{x|y} = \frac{\rho[xy]}{\rho[y]} = \frac{\rho[xy]}{\sum_i \rho[yi]}$$
(4.6)

so that all the pair transformation rates can be expressed in terms of pair densities and invasion rates. In equations (4.1) and (4.2) there are only three types of pair transformation, corresponding to the three kinds of edges in figure 4.1(c), examples of which were given in equations (4.3), (4.4), and (4.5). Use of the pair approximation simplifies the *q* terms so that the three generalised equations can be written as

$$R[ab \rightarrow aa] = \rho_{ab}r_a(1 + (Q - 1)q_{a|b})$$

$$R[aa \rightarrow ca] = \rho_{ca}r_c(Q - 1)q_{a|a}$$

$$R[ca \rightarrow ab] = \rho_{bc}r_b(Q - 1)q_{a|c}$$
(4.7)

for *a*, *b*, *c* as in equation (4.1). It might seem strange that equations (4.7) use the *unordered* density for the original pair (p_{aa} , p_{ab}), whereas *ordered* pair densities $\rho[xy]$ are used in the expansion (equation (4.6)) of the conditional terms $q_{c|a}$, $q_{a|b}$. This is because once the initial pair has been settled on, the ordering of the second pair relative to the first becomes important. In fact, because the pair in which the invasion occurs must always be a heterogeneous pair, it wouldn't make any difference if we used the ordered pair densities in equation (4.7) anyway.

Together, equations (4.1), (4.2), and (4.7) describe a system of six differential equations expressing the rates of change of all six unordered pair densities

in terms of the pair densities at the previous timestep and the three invasion rates.

4.3.5 Rate of change of singleton densities

Because every site has the same number of neighbours, the change in singleton densities can be found by summing the changes of the relevant pair densities. For example,

$$\frac{\mathrm{d}\rho_{\mathsf{r}}}{\mathrm{d}t} = \frac{\mathrm{d}\rho_{\mathsf{rr}}}{\mathrm{d}t} + \frac{1}{2} \cdot \frac{\mathrm{d}\rho_{\mathsf{rs}}}{\mathrm{d}t} + \frac{1}{2} \cdot \frac{\mathrm{d}\rho_{\mathsf{pr}}}{\mathrm{d}t}$$

The fraction $\frac{1}{2}$ occurs for the rates of change of non-homogeneous pairs because these pairs are unordered: ρ_{rs} is the density of both [rs] and [sr] pairs. The singleton densities are just the sum of ordered pairs, for example, ρ_r is calculated by summing over adjacent pairs of sites which have an r in the left hand position only. This sum can be expanded using (4.1) and (4.7), and becomes

$$\begin{split} \frac{\mathrm{d}\rho_{\mathrm{r}}}{\mathrm{d}t} &= R[\mathrm{rs} \rightarrow \mathrm{rr}] - R[\mathrm{rr} \rightarrow \mathrm{pr}] \\ &+ \frac{1}{2} \big(R[\mathrm{ss} \rightarrow \mathrm{rs}] + R[\mathrm{pr} \rightarrow \mathrm{rs}] - R[\mathrm{rs} \rightarrow \mathrm{rr}] - R[\mathrm{rs} \rightarrow \mathrm{sp}] \big) \\ &+ \frac{1}{2} \big(R[\mathrm{rr} \rightarrow \mathrm{pr}] + R[\mathrm{sp} \rightarrow \mathrm{pr}] - R[\mathrm{pr} \rightarrow \mathrm{pp}] - R[\mathrm{pr} \rightarrow \mathrm{rs}] \big) \\ &= \frac{1}{2} \big(R[\mathrm{rs} \rightarrow \mathrm{rr}] + R[\mathrm{ss} \rightarrow \mathrm{rs}] + R[\mathrm{sp} \rightarrow \mathrm{pr}] \\ &- R[\mathrm{rr} \rightarrow \mathrm{pr}] - R[\mathrm{rs} \rightarrow \mathrm{sp}] - R[\mathrm{pr} \rightarrow \mathrm{pp}] \big) \\ &= \frac{1}{2} \big(\rho_{\mathrm{rs}}r_{\mathrm{r}} \big(1 + (Q - 1)q_{\mathrm{r}|\mathrm{s}} \big) + \rho_{\mathrm{rs}}r_{\mathrm{r}} (Q - 1)q_{\mathrm{s}|\mathrm{s}} + \rho_{\mathrm{rs}}r_{\mathrm{r}} (Q - 1)q_{\mathrm{p}|\mathrm{s}} \\ &- \rho_{\mathrm{pr}}r_{\mathrm{p}} \big(Q - 1 \big) q_{\mathrm{r}|\mathrm{r}} - \rho_{\mathrm{pr}}r_{\mathrm{p}} \big(Q - 1 \big) q_{\mathrm{s}|\mathrm{r}} - \rho_{\mathrm{pr}}r_{\mathrm{p}} \big(1 + (Q - 1)(q_{\mathrm{r}|\mathrm{s}} + q_{\mathrm{s}|\mathrm{s}} + q_{\mathrm{p}|\mathrm{s}}) \big) \\ &- \rho_{\mathrm{pr}}r_{\mathrm{p}} \big(1 + (Q - 1)(q_{\mathrm{r}|\mathrm{r}} + q_{\mathrm{s}|\mathrm{r}} + q_{\mathrm{p}|\mathrm{r}}) \big) \Big). \end{split}$$

Because $q_{\mathbf{r}|x} + q_{\mathbf{s}|x} + q_{\mathbf{p}|x} = 1$, this becomes

$$\frac{\mathrm{d}\rho_{\mathsf{r}}}{\mathrm{d}t} = \frac{Q}{2}(\rho_{\mathsf{rs}}r_{\mathsf{r}} - \rho_{\mathsf{pr}}r_{\mathsf{p}}). \tag{4.8}$$

Equation (4.8) does not need to be part of the dynamical system but it is interesting to compare it with the mean field equation (1.1) on page 3, which has the same form except that pair densities (such as ρ_{rs}) are used in place



Figure 4.2: Examples of trajectories of species densities for the RPS pair approximation equations given in section 4.3. The black dot shows the initial singleton densities, the black line the trajectory of the densities over time, and the yellow dot the final stable densities. The red cross and dotted line show the mean field fixed point and trajectory for the same combination of initial densities and invasion rates. Initial densities are equal on the left and (0.4,0.2,0.4) on the right; invasion rates are (1,0.2,0.1) on the left and (1,0.5,0.6) on the right.

of the multiples of singleton densities (such as $\rho_r \rho_s$). The factor Q/2 appears because there are Q/2 times as many pairs as singletons for the reasons mentioned above in section 4.2.3.

4.4 Failure of the pair approximation for RPS

RPS depends on spatial interactions for stability, so one would expect the pair approximation, which includes more spatial information than the mean field

model, to be more likely to predict a stable outcome. But this is not the case. Sometimes the pair approximation equations predict a stable community at the same fixed point as the mean field, as shown on the left hand side of figure 4.2. More commonly, as shown on the right hand side of figure 4.2, the pair approximation shows ever increasing fluctuations around the fixed point until one of the two species goes extinct, in a similar way to the non-spatial, finite-population trajectory shown in figure 1.3 on page 6.

4.4.1 Importance of neighbourhood size

60

The first of the three equations in (4.7) suggests that the effect of the neighbourhood size Q is to change the relative frequency of the creation of homogeneous and heterogeneous pairs. For small Q, the equations suggest that homogeneous pairs will be produced more often than they are destroyed, other things being equal.

Exactly what the ultimate effect will be on the system's behaviour is difficult to predict, but for the examples in figure 4.2, *Q* appears to have an effect on the speed with which the system reaches its equilibrium. Convergence is faster when *Q* is low (figure 4.2, top diagrams), and slower when *Q* is high (bottom diagrams).

4.4.2 Extinctions more common than coexistence

The instability of the RPS pair approximation has been noted in Szabó et al. (2004), however they do not mention that there are cases in which the pair approximation correctly predicts stable coexistence. This is possibly because such cases are relatively rare, as shown in figure 4.3.

The small circles in the diagram indicate extinctions, and this is the predominant outcome for the five initial points shown by black crosses. Coexistence is only predicted when one of the three species has a much higher invasion rate than the other two (points in the corners of the triangle). Figure 4.2 suggests that where the pair approximation predicts coexistence, it predicts the same mixture of species as the mean field, and this appears to be the case, as the large circles in the first five (pair approximation) diagrams have the same shading as the corresponding circles in the mean field diagram.

As the initial densities are moved away from the centre, extinctions become more likely in the corners that are furthest from the initial point.


Figure 4.3: Equilibrium species densities for the RPS pair approximation system with Q = 4. Each circle represents a final species density. Small circles indicate extinctions with the winning species indicated by the colour of the circle (where rock = red, scissors = green, paper = blue). The larger circles indicate coexistence at equilibrium where the colour is a mixture with red, green, and blue proportional to the densities of rock, scissors and paper respectively (white circles are points without a result because they took too long to converge). The first five diagrams show the final pair approximation densities for five different initial densities indicated by black crosses. Each circle shows the density for a different set of invasion rates with the circle's position in the simplex determined by the mean field fixed point for those invasion rates. The mean field fixed point densities are shown in the sixth diagram.



Figure 4.4: Examples of the pair approximation trajectories with Q = 4 when a correlation quantified by ε is added to the initial pair densities. The resulting trajectories initially deviate slightly from those with uncorrelated initial pair densities in figure 4.2. The final densities are unaffected by the initial correlations.

4.4.3 Effects of initial pair densities

In all the examples given so far in figures 4.2 and 4.3, initial pair densities were uncorrelated, so that at time zero, $\rho_{xx} = \rho_x \rho_x$, and $\rho_{xy} = 2\rho_x \rho_y$ for $x \neq y$.

Figure 4.4 shows that the addition of correlations between neighbouring sites at the start of the simulation has little effect on the outcome. This is done by defining a small bias ε , which is the amount by which there are fewer than the expected uncorrelated number of heterogeneous pairs, and more than the expected number of homogeneous pairs, so that

$$\begin{split} \rho_{\rm rs} + \rho_{\rm sp} + \rho_{\rm pr} &= (1 - \varepsilon)(2\rho_{\rm r}\rho_{\rm s} + 2\rho_{\rm s}\rho_{\rm p} + 2\rho_{\rm p}\rho_{\rm r}),\\ \rho_{\rm rr} + \rho_{\rm ss} + \rho_{\rm pp} &= \rho_{\rm r}^2 + \rho_{\rm s}^2 + \rho_{\rm p}^2 + \varepsilon(2\rho_{\rm r}\rho_{\rm s} + 2\rho_{\rm s}\rho_{\rm p} + 2\rho_{\rm p}\rho_{\rm r}) \end{split}$$

The additional $\varepsilon(2\rho_r\rho_s + 2\rho_s\rho_p + 2\rho_p\rho_r)$ pairs are apportioned up among the

homogeneous pairs proportionally to what their densities would have been had there been no pair correlations. For example, if ρ'_{rr} and ρ'_{rs} are the densities of rr and rs pairs in the correlated case,

$$\begin{split} \rho_{\rm rr}' &= \rho_{\rm r}^2 (1 + \frac{\varepsilon (2\rho_{\rm r}\rho_{\rm s} + 2\rho_{\rm s}\rho_{\rm p} + 2\rho_{\rm p}\rho_{\rm r})}{\rho_{\rm r}^2 + \rho_{\rm s}^2 + \rho_{\rm p}^2}),\\ \rho_{\rm rs}' &= 2\rho_{\rm r}\rho_{\rm s}(1 - \varepsilon). \end{split}$$

The examples in figure 4.4 show that there is a perceptible perturbation to the initial trajectories of the system when ε is initially set to 0.5, but it has no effect on the final outcome.

4.4.4 Pair approximation compared to SCA simulation

How well does the pair approximation match the SCA simulation in those situations in which coexistence is predicted? Figure 4.3 suggests that the community composition predicted by the pair approximation is the same as that predicted by the mean field fixed point. Figure 4.5(a) shows that this prediction can be more extreme than the result obtained in the SCA simulation. This is especially true when one of the three invasion rates is high relative to the other two.

Figure 4.5(b) and (c) show how the pair correlations change over time according to the two models. In the pair approximation, the proportion of homogeneous pairs climbs steeply and then settles into oscillations which slowly settle down and stabilise. In the SCA simulation, homogeneous pairs rise more smoothly, and barely oscillate once they are close to their equilibrium level. The pair approximation significantly underestimates the number of homogeneous pairs, which should be expected considering that a lot of information about local correlations is thrown away by the pair approximation assumption given in equation (4.6).

4.4.5 Summary

In general, the pair approximation fails to predict the essential feature of the spatial RPS system: large regions of stable coexistence of all three species. In those cases where the pair approximation gives good results for the equilibrium species densities, it significantly underestimates the amount of local correlation.



Figure 4.5: Comparison of the pair approximation to the SCA simulation. (a) Trajectories of the mean field (dotted red line), pair approximation (grey line) and 700×700 SCA simulation (blue line), for equal initial densities and invasion rates (1,0.2,0.25). (b) and (c) show the pair approximation's predicted proportions of homo- and heterogeneous pairs (grey solid and dashed lines) with the actual proportions of homo- and heterogeneous pairs counted during the SCA simulation over time (blue solid and dashed lines).

4.5 2×2 approximations

Hiebeler (1997) has compared the accuracy of mean field and pair approximation models with approximations that model the dynamics of larger (2 \times 2 and 4 \times 1) local blocks of cells, using a simple one-species basic contact process like that modelled by Durrett and Levin (1994b). He found that a model using approximations based on 2 \times 2 blocks of cells will predict patch occupancy probabilities about twice as accurately as the 2 \times 1 pair approximation



Figure 4.6: The rock-paper-scissors system: dynamics of 2×2 blocks of sites.

model, at the cost of significant complexity in the model. He also found the 2 \times 2 approximation model to be more accurate than the 4 \times 1.

A similar effect happens with the RPS system: a local structure approximation based on 2×2 blocks of cells turns out to be much better at predicting the stability of the system than the pair approximation, although it is not successful at predicting the exact equilibrium species densities.

4.5.1 **RPS and 2** \times 2 blocks

To apply the 2 × 2 approximation to the RPS ecosystem, we need to work out all the possible states of a 2 × 2 block of cells. Fortunately, it's not necessary to keep track of all possible combinations of three states in each of four positions, because there are $3^4 = 81$ blocks of the form $\begin{bmatrix} wx \\ yz \end{bmatrix}$ where $w, x, y, z \in \{r, s, p\}$. Some blocks are just rotations of others, and we can assume that they will occur with the same frequency because the cellular automata rules are symmetric. So if $\rho \begin{bmatrix} wx \\ yz \end{bmatrix}$ is the proportion of $\begin{bmatrix} wx \\ yz \end{bmatrix}$ blocks in the grid, then

$$\rho[_{yz}^{wx}] = \rho[_{yx}^{zw}] = \rho[_{xw}^{yz}] = \rho[_{wz}^{xy}], \ w, x, y, z \in \{\mathsf{r}, \mathsf{s}, \mathsf{p}\}.$$
(4.9)

Similarly, the two blocks $\begin{bmatrix} xx \\ yz \end{bmatrix}$ and $\begin{bmatrix} xx \\ zy \end{bmatrix}$ will occur with the same frequency if we are using what Durrett and Levin (1994b) call 'von Neumann neighbourhoods', where each cell can only interact with its closest four neighbours on the grid, or 'Moore neighbourhoods', where each cell has eight neighbours, including the four diagonal neighbours. This is because for each cell in the first block, there is a corresponding cell in the second block in the same state, and with the same immediate neighbours. I will therefore also assume that

$$\rho[_{yz}^{xx}] = \rho[_{zy}^{xx}], \ x, y, z \in \{\mathsf{r}, \mathsf{s}, \mathsf{p}\}.$$
(4.10)

With Moore neighbourhoods, it would also be possible to assume that $\begin{bmatrix} xx \\ zy \end{bmatrix}$ and $\begin{bmatrix} xy \\ zx \end{bmatrix}$ occur with the same frequency, but this will not be true in general with von Neumann neighbourhoods, which is the kind of neighbourhood structure I will assume for the remainder of this section.

The equalities in (4.9) and (4.10) allow the reduction of the 81 ordered blocks to 21 unordered block types. These 21 states are shown figure 4.6, with all the possible transitions between them when an invasion occurs at one site only.

4.5.2 Block densities, ordered and unordered

Using the same notation as for the pairs, I have referred to blocks in which the order of sites, rotation and reflection is important with square bracketed terms like $\begin{bmatrix} wx \\ yz \end{bmatrix}$. For the block *types* in which the order is unimportant, I omit the square brackets and use terms like $\begin{bmatrix} wx \\ yz \end{bmatrix}$. As with pairs, $\rho \begin{bmatrix} wx \\ yz \end{bmatrix}$ and $\rho \begin{bmatrix} wx \\ yz \end{bmatrix}$ are used for the densities of ordered and unordered blocks.

The block densities are simpler to calculate than they were for the pairs because 2 × 2 blocks can be counted so that there are the same number of blocks, N, as there are single sites. By contrast, in the pair approximation there were NQ/2 pairs for N sites. The difference is because of the need (for example when Q = 4) to count all north-south as well as east-west pairs to cover all the pairs in the grid, whereas it is possible to cover all the 2 × 2 blocks with a simpler tiling over the grid.

It will still be necessary to make use of the fact that each unordered block type may count more than one ordered block. For example, equation (4.9) implies that

$$\rho_{\rm rr}^{\rm pr} = \rho[{\rm rr}^{\rm pr}] + \rho[{\rm rr}^{\rm rp}] + \rho[{\rm rr}^{\rm rr}] + \rho[{\rm rr}^{\rm rr}],$$

and equations (4.9) and (4.10) together imply that

$$\rho_{\rm sr}^{\rm pr} = \rho[{\rm sr}^{\rm pr}] + \rho[{\rm sp}^{\rm sp}] + \rho[{\rm rs}^{\rm rs}] + \rho[{\rm rs}^{\rm rr}] + \rho[{\rm sr}^{\rm sr}] + \rho[{\rm rs}^{\rm rs}] + \rho[{\rm rs}^{\rm rs}] + \rho[{\rm rs}^{\rm rr}] + \rho[{\rm rs}^{\rm$$

4.5. 2×2 APPROXIMATIONS

The complete relationship between ordered and unordered blocks can be characterised by the relation *ordered*:

$$ordered(^{aa}_{aa}) = \{ [^{aa}_{aa}] \}$$

$$ordered(^{ca}_{aa}) = \{ [^{ca}_{aa}], [^{ac}_{aa}], [^{aa}_{ac}], [^{aa}_{ca}] \}$$

$$ordered(^{aa}_{ba}) = \{ [^{aa}_{ba}], [^{ba}_{aa}], [^{ab}_{aa}], [^{aa}_{ab}] \}$$

$$ordered(^{cc}_{aa}) = \{ [^{cc}_{aa}], [^{ac}_{ac}], [^{ca}_{cc}], [^{ca}_{ca}] \}$$

$$ordered(^{ab}_{ba}) = \{ [^{ab}_{ba}], [^{ba}_{ab}] \}$$

$$ordered(^{ca}_{ab}) = \{ [^{ca}_{ab}], [^{ac}_{ba}], [^{ab}_{ca}], [^{ab}_{ca}] \}$$

$$ordered(^{ca}_{ab}) = \{ [^{ca}_{ab}], [^{ac}_{ba}], [^{ab}_{ca}], [^{ab}_{ca}] \}$$

$$ordered(^{ca}_{ba}) = \{ [^{ca}_{ba}], [^{ac}_{ba}], [^{ab}_{ca}], [^{ab}_{ca}], [^{cb}_{ca}], [^{ac}_{ab}], [^{ac}_{ab}] \}$$

for *a*, *b*, *c* as in equation (4.1). I will call the inverse relation *unordered*(), a function I use below when deriving the rates of change of all the blocks.

4.5.3 Rate of change of block types

The rate of change of any type of 2×2 block can be determined in the same way as for pairs of cells; it is the rate at which blocks of that type are created minus the rate at which blocks of that type are destroyed. The equations can be written by looking at figure 4.6, where each node of the diagram represents one type of 2×2 block. There is one equation for each node in figure 4.6, and each equation will have one term for each edge which is connected to the corresponding node. The incoming edges on the diagram become positive terms in the equation, and the outgoing edges become negative terms in the equation.

Also, the structure of the upper seven nodes in figure 4.6 is identical to the structure of the seven nodes at the bottom right of the graph and the seven nodes at the bottom left of the graph, so the 21 equations can be summarised by the seven below, where $a \in \{r, s, p\}, b = prey(a), c = predator(a)$.

$$\begin{aligned} \frac{\mathrm{d}\rho_{aa}^{aa}}{\mathrm{d}t} &= R[_{ab}^{aa} \to _{aa}^{aa}] - R[_{aa}^{aa} \to _{aa}^{ca}] \\ \frac{\mathrm{d}\rho_{aa}^{ca}}{\mathrm{d}t} &= R[_{aa}^{aa} \to _{aa}^{ca}] + R[_{ba}^{ca} \to _{aa}^{ca}] + R[_{ab}^{ca} \to _{aa}^{ca}] \\ &- R[_{aa}^{ca} \to _{ac}^{ca}] - R[_{aa}^{ca} \to _{aa}^{cc}] - R[_{aa}^{ca} \to _{aa}^{ca}] \\ \frac{\mathrm{d}\rho_{aa}}{\mathrm{d}t} &= R[_{aa}^{ca} \to _{ab}^{aa}] + R[_{ab}^{ab} \to _{ab}^{aa}] + R[_{bb}^{ab} \to _{ab}^{aa}] \\ &- R[_{aa}^{ca} \to _{ab}^{aa}] + R[_{ab}^{ab} \to _{ab}^{aa}] + R[_{bb}^{aa} \to _{ab}^{ca}] \\ &- R[_{aa}^{aa} \to _{aa}^{aa}] - R[_{ab}^{aa} \to _{ba}^{ca}] - R[_{ab}^{aa} \to _{ab}^{ca}] \\ &- R[_{aa}^{aa} \to _{aa}^{aa}] - R[_{ab}^{aa} \to _{ba}^{ca}] - R[_{aa}^{aa} \to _{ab}^{ca}] \end{aligned}$$

$$\frac{d\rho_{ab}}{dt} = R[^{ca}_{ba} \rightarrow ^{ab}_{ab}] + R[^{ab}_{bb} \rightarrow ^{ab}_{ab}] - R[^{ab}_{ab} \rightarrow ^{aa}_{ab}] - R[^{ab}_{ab} \rightarrow ^{ab}_{cb}]$$

$$\frac{d\rho_{ab}}{dt} = R[^{aa}_{ab} \rightarrow ^{ca}_{ba}] + R[^{ca}_{ca} \rightarrow ^{ca}_{ba}] + R[^{bb}_{bc} \rightarrow ^{ca}_{ba}]$$

$$- R[^{ca}_{ba} \rightarrow ^{ab}_{ab}] - R[^{ca}_{ba} \rightarrow ^{ca}_{ca}] - R[^{ca}_{ba} \rightarrow ^{ca}_{ca}]$$

$$\frac{d\rho_{ba}}{dt} = R[^{aa}_{ab} \rightarrow ^{ca}_{ab}] + R[^{cc}_{aa} \rightarrow ^{ca}_{ab}] + R[^{bb}_{bc} \rightarrow ^{ca}_{ca}]$$

$$- R[^{ca}_{ab} \rightarrow ^{ca}_{ab}] + R[^{cc}_{aa} \rightarrow ^{ca}_{ab}] + R[^{ab}_{bc} \rightarrow ^{ca}_{ab}] + R[^{ab}_{cb} \rightarrow ^{ca}_{ab}]$$

$$- R[^{ca}_{ab} \rightarrow ^{ca}_{aa}] - R[^{ca}_{ab} \rightarrow ^{bc}_{ca}] - R[^{ca}_{ab} \rightarrow ^{bc}_{ac}] - R[^{ca}_{ab} \rightarrow ^{bc}_{ab}]$$

$$(4.11)$$

4.5.4 Block transformation rates

The $R[_{uv}^{st} \rightarrow \frac{wx}{yz}]$ terms can be translated using invasion rates and conditional probabilities similar to the conditional probabilities of the form $q_{x|yz}$ that were used in the pair approximations.

It is important to remember that all the basic events in the system are invasions from a single cell to its neighbour, so these basic events depend on the densities of pairs, not blocks. In other words, although the events are all what Rand (1999) calls 'edge events', the effects of those edge events must be described in terms of blocks, rather than edges.

Example: $R[_{rr}^{rr} \rightarrow _{rr}^{pr}]$

For example, consider the rate at which r_{rr}^{rr} blocks are transformed into r_{rr}^{pr} blocks, $R[r_{rr}^{rr} \rightarrow r_{rr}^{pr}]$. This can only happen when there is an existing r_{rr}^{rr} block which is successfully invaded by a neighbouring p cell, so the only event which is important is the invasion $p \rightarrow r$, an event which occurs in the system at the overall rate of $r_p \rho_{pr}$.

When the $p \rightarrow r$ event takes place, the number of r_r^{pr} blocks that are created is determined by the probability that the r of the original pr pair in which the invasion takes place was part of a r_r^{rr} block, shown in the following picture:

p—	→r	r
	r	r

The transformed block is on the right, and the overlapping pair in which the $p \rightarrow r$ invasion takes place is shown on the left. The chance that a $_{rr}^{pr}$ block is created by the invasion is described by $q_{r'/rp}$, the conditional probability that all the other sites in the overlapping r's block are also in state r, given the

4.5. 2×2 APPROXIMATIONS

original invasion pair.

The r in the invasion pair is potentially inside two transformed blocks: the lower one, like the transformed block in the diagram, and an upper one as well, so the overall transformation rate has a factor of two and is described by

$$R[_{\rm rr}^{\rm rr} \rightarrow _{\rm rr}^{\rm pr}] = 2r_{\rm p}\rho_{\rm pr}q_{\rm r''/rp}.$$

Example: $R[_{sr}^{rr} \rightarrow _{rr}^{rr}]$

A slightly more complicated example is the rate at which r_{rr}^{rr} blocks are created from s_{rr}^{rr} blocks, $R[s_{rr}^{rr} \rightarrow r_{rr}^{rr}]$. It is more complicated because of the possibility that the invasion pair is itself part of a transformed block.

In this case the only relevant invasion is $r \rightarrow s$, after which the invasion pair will be in state rr, so if the immediate neighbours of the invasion pair were also in state rr before the invasion, this will create a new r_{rr}^{rr} block. The relevant conditional probability is $q_{r|s}$, the probability that two neighbours on the same side of an existing rs pair are in state rr. These pairs are in addition to any others that overlapped the original s without overlapping the original r, and so the overall rate is

$$R[{}_{\rm sr}^{\rm rr} \rightarrow {}_{\rm rr}^{\rm rr}] = 2r_{\rm r}\rho_{\rm rs}(q_{\rm r|s}{}_{\rm r|s}{} + q_{\rm rr}{}_{\rm rsr}{}).$$

Unique block transformations

The 54 arrows shown in figure 4.6 are made up of three sets of 18 arrows, one set for each species. Only the 18 unique types of block transformation rate are listed here:

$$\begin{split} &R[\overset{aa}{aa} \rightarrow \overset{ca}{aa}] = 2r_{c}\rho_{ca}q^{aa}_{a'ac} \\ &R[\overset{ca}{aa} \rightarrow \overset{aa}{ba}] = 2r_{b}\rho_{bc}q^{aa}_{a'cb} \\ &R[\overset{ca}{aa} \rightarrow \overset{cc}{aa}] = 2r_{c}\rho_{ca}(q^{a}_{a|a} + q^{aa}_{c'ac} + q^{ac}_{a'ac}) \\ &R[\overset{ca}{aa} \rightarrow \overset{ca}{ac}] = 2r_{c}\rho_{ca}q^{ca}_{a'ac} \\ &R[\overset{aa}{ba} \rightarrow \overset{ca}{aa}] = 2r_{a}\rho_{ab}(q^{a}_{a|b} + q^{a}_{a'ba}) \\ &R[\overset{aa}{ba} \rightarrow \overset{ca}{ab}] = 2r_{c}\rho_{ca}q^{ba}_{a'ac} \\ &R[\overset{aa}{ba} \rightarrow \overset{ca}{ab}] = 2r_{c}\rho_{ca}(q^{a}_{ab} + q^{a}_{a'ba}) \\ &R[\overset{aa}{ba} \rightarrow \overset{ca}{ab}] = 2r_{c}\rho_{ca}(q^{a}_{ab} + q^{a}_{b'ac}) \\ &R[\overset{aa}{ba} \rightarrow \overset{ca}{ba}] = 2r_{c}\rho_{ca}(q^{c}_{ab} + q^{a}_{b'ac}) \\ &R[\overset{cc}{aa} \rightarrow \overset{cc}{ac}] = 2r_{c}\rho_{ca}(q^{c}_{a|a} + q^{a}_{a'cc} + q^{c}_{c'ac}) \\ &R[\overset{cc}{aa} \rightarrow \overset{ca}{ba}] = 2r_{b}\rho_{bc}(q^{c}_{a} + q^{a}_{a'cb}) \\ &R[\overset{ab}{ba} \rightarrow \overset{aa}{ba}] = 2r_{a}\rho_{ab}(q^{b}_{a|b} + q^{b}_{a'ba}) \end{split}$$

$$\begin{aligned} R[^{ab}_{ba} \rightarrow ^{ab}_{bc}] &= 2r_c \rho_{ca} q_{\frac{ab}{b}'ac} \\ R[^{ca}_{ab} \rightarrow ^{ca}_{aa}] &= 2r_a \rho_{ab} (q_{\frac{c}{a}|b} + q_{\frac{c}{a}'ba}) \\ R[^{ca}_{ab} \rightarrow ^{ab}_{ba}] &= 2r_b \rho_{bc} q_{\frac{b}{a}'cb} \\ R[^{ca}_{ab} \rightarrow ^{bc}_{ac}] &= 2r_c \rho_{ca} (q_{\frac{b}{a}|a}^{ac} + q_{\frac{a}{b}'ac} + q_{\frac{b}{a}'cb}) \\ R[^{ca}_{ba} \rightarrow ^{ca}_{aa}] &= 2r_a \rho_{ab} (q_{\frac{a}{a}|b}^{aa} + q_{\frac{a}{a}'ba} + q_{\frac{a}{a}'ba}) \\ R[^{ca}_{ba} \rightarrow ^{ca}_{aa}] &= 2r_b \rho_{bc} (q_{\frac{a}{a}|b}^{ab} + q_{\frac{a}{a}'ba} + q_{\frac{a}{a}'ba}) \\ R[^{ca}_{ba} \rightarrow ^{ca}_{bb}] &= 2r_b \rho_{bc} (q_{\frac{a}{a}|b}^{ab} + q_{\frac{a}{a}'ba}^{ad} + q_{\frac{a}{a}'ba}) \\ R[^{ca}_{ba} \rightarrow ^{bc}_{ca}] &= 2r_c \rho_{ca} (q_{\frac{b}{b}'ac}^{ca} + q_{\frac{b}{a}'ac}^{cb}) \\ R[^{ca}_{ba} \rightarrow ^{bc}_{ac}] &= 2r_c \rho_{ca} (q_{\frac{b}{b}'ac}^{ab} + q_{\frac{a}{a}'ba}^{cb}) \end{aligned}$$

$$(4.12)$$

4.5.5 The '2 \times 2 block' approximation

The conditional probabilities $q_{\frac{vw}{x}/y_z}$ and $q_{\frac{w}{y}|z}$ terms used in equations (4.12) must now be described using only 2 × 2 block densities.

 $q_{y|z}^{w|x}$ requires no simplifying approximation, and can be defined using the same notion of conditional probabilities as used for the pair approximation:

$$q_{\substack{w|x\\y|z}} = \frac{\rho[\substack{wx\\yz}]}{\rho[\substack{x\\z}]} = \frac{\rho[\substack{wx\\yz}]}{\sum_{i,j}\rho[\substack{ix\\jz}]}$$

On the other hand the $q_{x'y_z}^{vu}$ terms really depend on the densities of blocks of six cells,

$$q_{x}^{vw}_{yz} = rac{
ho[vw*]}{
ho[yz]},$$

where the * in the top line indicates a cell whose state doesn't matter.

Probabilities of triplets were approximated in terms of pair densities using the pair approximation given in section 4.3.4. The analogous local structure approximation needed here is the assumption that in blocks of six such as $\begin{bmatrix} uvw \\ xyz \end{bmatrix}$, the $\begin{bmatrix} u \\ x \end{bmatrix}$ part is independent of the $\begin{bmatrix} w \\ z \end{bmatrix}$ part, given the $\begin{bmatrix} v \\ y \end{bmatrix}$ in the middle:

$$q_{\substack{u \mid vw \\ x \mid yz}} \approx q_{\substack{u \mid v \\ x \mid y}}.$$

This assumption allows the conditional probabilities to be rewritten using 2 \times 2 blocks:

$$q_{xyyz} = \frac{\rho[_{xyz}^{vw*}]}{\rho[yz]} \approx \frac{\rho[_{xy}^{vw}]\rho[_{yz}^{w*}]}{\rho[_{y}^{w}]\rho[yz]} = \frac{\rho[_{xy}^{vw}]\sum_{i}\rho[_{yz}^{wi}]}{\sum_{i,j}\rho[_{yj}^{wi}]\sum_{i,j}\rho[_{yz}^{ij}]}.$$
(4.13)



Figure 4.7: Comparison of the 2 \times 2 approximation and mean field model for the same initial densities (shown by the black dot) and invasion rates as those in figure 4.2. The red dotted line shows the stable orbit of the three species densities under the mean field model, and the red cross shows the mean field fixed point. The green line and green point show the trajectory and rest point for the 2 \times 2 model. In both cases the initial densities of the 2 \times 2 blocks were chosen to represent a situation in which the states of all sites are uncorrelated.

The summation over blocks, and the '2 \times 2 block' approximation close off the equations in (4.11) and (4.12) so that they can all be defined only in terms of 2 \times 2 block densities and the invasion rates.

And because it is more convenient to keep track only of the densities of the unordered blocks, the *unordered()* function mentioned previously can be used to get the densities of the ordered blocks from equation (4.13) in terms of unordered block densities, by taking the density of the corresponding unordered block and dividing by the number of ordered blocks represented by that unordered block:

$$\rho[_{yz}^{wx}] = \frac{\rho_{unordered}([_{yz}^{wx}])}{\# ordered((unordered([_{yz}^{wx}])))}.$$

4.6 Behaviour of the 2×2 model

Unlike the pair approximation, the 2×2 block approximation correctly picks that the spatial RPS system tends towards a stable attractor in which all three species are present. However, it fails to accurately predict the composition of the stable community, and for most parameter values the mean field fixed point gives a better approximation.

4.6.1 Stability

Figure 4.7 shows that the three species densities tend to settle to a stable fixed point under the 2×2 approximation model. In this respect it reproduces the result of the explicit SCA simulation in which there is a stable point attractor. By contrast, the mean field model predicted a stable orbit but not a stable point, and under the pair approximation, the central fixed point is a repeller for most parameter values.

In the cases where the pair approximation model gives a stable central point, such as in the left hand diagram of figure 4.2 on page 59, the 2 \times 2 approximation system converges to the fixed point more quickly, with fewer cycles (figure 4.7).

4.6.2 Ecosystem composition

Figure 4.7 also clearly shows that the stable point of the 2×2 approximation is in a different place from the fixed point of the mean field model. The 2×2 approximation predicts that the prey of the fastest invader will end up with a lower density (at the expense of the other two species) compared with the mean field result.

A variety of central fixed points are shown in figure 4.8 for the mean field, pair approximation, 2×2 approximation, and SCA simulation. Each of the 15 colours represents a set of invasion rates, and each different shaped symbols shows the central fixed point (if it exists) for one of the four models. The accuracy of the approximations appears to depend on the relative invasion rates of the three species in the following ways:

- (1) When the invasion rates are evenly matched, the mean field, 2×2 approximation, and simulation all give similar results.
- (2) When one species grows much slower than the other two (at the edges of the diagram) the mean field fixed point stays close to the simulation result, but the 2×2 approximation underestimates the density of the slow-growing species, and overestimates the density of its prey.
- (3) When one species grows much faster than the other two (in the corners of the diagram), the mean field does better than the 2×2 but both approximations predict that the faster-growing (and lowest density) species is more vulnerable than it actually is according to the spatial simulation.



Figure 4.8: Comparison of fixed point locations of the species densities for 15 different values of the invasion rates under the mean field, pair and 2×2 approximation models, and SCA simulation. Each colour represents one set of invasion rates. The diamond, circle, and square symbols show fixed points for the mean field, pair approximation, and 2×2 approximation respectively. The densities under the simulations continue to fluctuate indefinitely. The plotted lines show the species densities for the last 30% of the simulation, and the stars show the densities at the termination of the simulation. Stable points for the pair approximation are only shown for the three cases in which extinctions are not predicted by that model.

The difference between each approximation and a set of simulation runs is shown in figure 4.9. The pair approximation, which predicts species extinctions in most of the parameter space, gives by far the worst result, and the other two are more accurate. But the mean field gives a slightly better result than the 2 \times 2 approximation everywhere. So despite their more accurate treatment of space, neither the pair nor the 2 \times 2 approximation model accurately predicts the composition of species in the RPS system.



Figure 4.9: Difference in predicted species densities between each of the three approximations and the result of spatial SCA simulations on a 700 × 700 grid, for a range of invasion rates. The error measurement used in each case, and shown on the colour scale, is the distance in species-density space between the simulation and the approximation, $\sqrt{\sum_{i \in \{r,s,p\}} (\rho_i^{sim} - \rho_i^{approx})^2}$. The ρ_i^{sim} values used here are the mean species densities for the last 30% of timesteps in a 600-generation simulation, equivalent to the mean of all the points on the lines plotted for the simulation in figure 4.8.

4.6.3 Neighbour correlations and species clustering

While the pair and 2×2 approximations are poor at finding the RPS equilibrium species densities, they are an improvement on the mean field approximation when predicting correlations between the states of neighbouring grid sites.

Correlations between neighbouring pairs are shown for a simulation run and for all three approximations in figure 4.10 for three sets of invasion rates. The



Figure 4.10: Predictions of densities of cell neighbours for three different sets of growth rates by simulation, 2×2 approximation, pair approximation and mean field techniques. The first row shows the results of explicit spatial simulations of the RPS ecosystem using three different sets of growth rates, on a 700×700 grid after 800 time steps. The red, green and blue areas represent sites occupied respectively by rocks, scissors, and paper. The second row shows, for the three sets of growth rates, the state of all the cells neighbouring rocks, firstly from the result of the simulation (Sim), and the same proportions estimated using the pair approximation (PA), 2×2 approximation (2x2), and mean field (MF) (at its fixed point). The third and fourth rows show the expected proportions of cells neighbouring scissors and paper respectively. In (b), some pair approximation data are missing where the pair approximation predicts the extinction of rocks and scissors. The mathematical approximations underestimate the extent of clustering in the simulation.

top row shows the state of a simulation, and the phenomenon that clustering is more pronounced when invasion rates are mismatched. The second, third and fourth rows show the expected proportions of neighbours of cells occupied by each species.

When one species grows much faster than the other two ($r_r = 0.2$, $r_s = 1$, $r_p = 0.2$, figure 4.10(a)), the grid becomes very clustered in the simulation, but is less clustered as the growth rates become more equal in the diagrams in (b) and (c). All three mathematical approximations underestimate the correlations between neighbouring sites, and they tend to make worse predictions with larger amounts of clustering.

The local structure approximations tend to do better than the mean field, but this is not always the case. In (a) and (c), the pair approximation does best, followed by the 2×2 and then the mean field. In (b), the pair approximation predicts an extinction, and the 2×2 approximation predicts the same-species correlations better than the mean field, but the mean field is actually slightly more accurate than the 2×2 in two of the nine correlations, the scissors-rock and paper-rock cases.

Clustering and inaccuracy in densities

Although the clustering is most extreme at the corners of the species density simplex, and although it is at these points that the species densities predicted by the approximations are least accurate, the underestimation of clustering in these regions cannot be the primary reason for the failure of the density predictions. This is because the mean field does better than the 2×2 at predicting the densities but is generally worse at predicting the clustering.

4.7 Discussion

76

The results presented in this chapter show that local structure approximations are not always an appropriate way of modelling spatial systems of interacting individuals. In RPS communities, the two local structure approximations models fail to capture the effects of a spatially structured population, and give misleading results (predicting extinctions rather than coexistence) in many cases.

More generally, this suggests that explicit agent-based spatial computer models will remain an important tool in theoretical ecology. At the very least they

4.7. DISCUSSION

will be needed to verify that infinite-population local structure approximations are able to correctly model the effects of heterogeneous spatial structure for a given model ecosystem.

Chapter 5

Direct approximation of the stabilising effect of spatial structure in the mean field

The local structure approximations discussed in the last chapter are poor approximations (compared with the mean field) when the goal is to predict equilibrium species densities in a spatial RPS game. But sometimes equilibrium densities are not the property of interest. For example, there are times when it would be useful to estimate the *vulnerability* of a RPS ecosystem to collapse, which can be represented by the minimum species density over time.

Given that local spatial interactions have a stabilising effect, it seems likely that the mean field model would be inadequate for this goal, because its oscillations are constant and do not exhibit any dampening down over time towards a stable point. To address this problem, in this section I describe a very simple approximation which is similar to the mean field equations but with an explicit spatial stability factor forcing the fixed point to be an attractor.

I show that the model does not help to predict the minimum species density of the simulation, because SCA simulations usually exhibit a large initial oscillation; the simulated ecosystem's vulnerability is higher on its first orbit than the mean field predicts. The initial instability of the SCA is due to the fact that sites are initialised in a random state.

5.1 Explicit spatial stability

Heterogeneous space has a stabilising effect on the RPS cellular ecosystem, and causes the system to be attracted to a particular combination of species densities, which is in most cases quite close to the fixed point of the mean field equations (1.1), on page 3.

An interesting question is whether a model that explicitly adds spatial stability to the mean field can be useful. Such a model would be much less computationally intensive than the SCA, would have the same fixed point as the mean field model, but the fixed point would be an attractor. The model should have more predictive accuracy than the mean field model, but at the cost of the mysterious 'space factor', which is just assumed rather than justified using the nature of the known spatial structure of the system.

The spatial stability factor *h*

I propose a single constant parameter h, the spatial stability factor, which represents the amount of attraction towards the fixed point in one generation.

The changes in densities of the three species $\{r, s, p\}$ using *h* are:

$$\dot{\rho_{\rm r}} = \rho_{\rm r} (\rho_{\rm s} r_{\rm r} - \rho_{\rm p} r_{\rm p}) (1 - h) - (\rho_{\rm r} - \frac{r_{\rm s}}{r_{\rm r} + r_{\rm s} + r_{\rm p}}) h$$

$$\dot{\rho_{\rm s}} = \rho_{\rm s} (\rho_{\rm p} r_{\rm s} - \rho_{\rm r} r_{\rm r}) (1 - h) - (\rho_{\rm s} - \frac{r_{\rm p}}{r_{\rm r} + r_{\rm s} + r_{\rm p}}) h$$

$$\dot{\rho_{\rm p}} = \rho_{\rm p} (\rho_{\rm r} r_{\rm p} - \rho_{\rm s} r_{\rm s}) (1 - h) - (\rho_{\rm p} - \frac{r_{\rm r}}{r_{\rm r} + r_{\rm s} + r_{\rm p}}) h$$
(5.1)

where the first terms are the same as in the mean field equation, multiplied by 1 - h, and the second terms change the densities explicitly in the direction of the mean field fixed point where $\rho_r = r_s/(r_r + r_s + r_p)$, $\rho_s = r_p/(r_r + r_s + r_p)$, $\rho_p = r_r/(r_r + r_s + r_p)$. Equations (5.1) have therefore deliberately been written so that they have a stable fixed point at the same place as the fixed point of the mean field equations.

5.2 Behaviour of the approximation

Figure 5.1 shows an example of the behaviour of equations (5.1) for the same initial densities and invasion rates for different values of h. When h = 0, the



Figure 5.1: Behaviour of equations (5.1) for four values of h, with initial densities (0.15,0.15,0.7) (marked by a star), and fixed invasion rates (0.25,0.3,0.45), which determine the fixed point (red dot).

system is the same as the mean field system in equations (1.1), but as h is increased the stability of the system is also increased.

5.2.1 Extinction risk

By its definition, this quantification of spatial stability is not useful in predicting equilibrium species densities than the mean field model, but such a system could provide a method for approximating the risk of an extinction in the system.

If one of the three species goes extinct, then the species which preys on it will go extinct soon afterwards. For a given set of invasion rates R and initial densities P, there is an extinction risk which can be described in terms of the minimum species density of any of the three species over time, or, on a simplex plot, by the smallest distance to an edge at which one of the densities is equal to zero.



Figure 5.2: Minimum species densities for a spatial SCA (500×500 grid, green), mean field approximation (red), and equations (5.1) (with h = 0.01, blue). Initial densities are (0.15,0.25,0.6) (black cross) and invasion rates (0.25,0.45,0.3). The stable point is marked with a star.

5.2.2 Predictions of extinction risk

Figure 5.2 shows a comparison between a spatial SCA simulation and the mean field and explicit stability approximations.

On its first orbit around the fixed point, the SCA simulation actually has lower minimum densities (and therefore higher predicted extinction risk) than the mean field model for all three species. The model with explicit stability is worse at predicting minimum densities than the mean field, even with a very low h = 0.01, although it shows the stable nature of the fixed point. This is all because the SCA simulation is very unstable on its first orbit.

The behaviour of the SCA simulation in figure 5.2 is fairly typical. Table 5.1 shows the average error of both approximations when predicting minimum species densities across hundreds of separate simulations, each with randomly picked invasion rates and initial densities. The mean field is a better predictor of minimum densities than equations (5.1) because the minimum densities always occur on the first orbit, but even the mean field picks minima which are consistently too high.

Table 5.1 shows that in the SCA, the greatest extinction risk happens when the second species reaches its minimum, because the second minimum is lower, on average, than either the first or third minima. It appears that the

	1 st min	2 nd min	3 rd min	mean
SCA Simulation	0.158	0.055	0.091	0.101
Mean field	0.165	0.072	0.114	0.117
(diff. from sim.)	+0.007	+0.017	+0.023	+0.016
Eqns (5.1), <i>h</i> = 0.01	0.182	0.094	0.166	0.147
(diff. from sim.)	+0.024	+0.038	+0.075	+0.046

Table 5.1: Difference between simulated minimum species densities (on a 500×500 SCA grid) and predicted minimum densities for mean field and explicit space models, based on the average over a total of 1204 simulations in which none of the three species densities fell below 0.0001 during the first 400 generations. Each simulation started with random species densities and invasion rates.

first species tends to reach its minimum density before the SCA reaches its maximum instability, and that the third species reaches its minimum after the system has begun to stabilise.

I suspect that there is no simple modification of the mean field equations that can capture the initial instability and subsequent stability of the spatial SCA simulation for RPS systems. A system that could capture the initial instability would have to be non-Markovian, because both early trajectories going outwards, and later trajectories heading towards the central fixed point can pass through the same point on the way. The SCA contains a kind of 'memory' effect in the configuration of its sites. The most obvious visible instantiation of this memory is the amount of clustering of same-species individuals. An SCA that is initialised at random tends to go through a stage in which very large clusters sweep across the grid before breaking up into smaller ones like those of figure 2.5 on page 21.

The initial non-Markovian behaviour of the SCA is a consequence of the unrealistic initialisation of the simulation in which each site is put into a random state independently of its neighbours. If the well-mixed state is judged to be an improbable starting point, then it may be unnecessary to initialise a model of a real ecosystem in this way. The approximation presented here may therefore be applicable to systems which begin with small perturbations from the equilibrium, or which begin with a near-monoculture of one species and so is already in its most 'unstable' state.

Part III

The evolution of competitive cycles

Chapter 6

The evolution of intransitive competition from intraspecific competition

This chapter investigates a possible solution to the problem of how stable cyclic competitive ecosystems such as RPS can originate. The proposed solution involves a SCA model of a two-species system with intraspecific competition, devised by Marcus Frean. In this model a speciation event separates one of the two species into two disjoint morphs. These two morphs, along with the monomorphic second species, make three separate groups which compete together in an intransitive cycle.

Section 6.1 provides an introduction to the problem and describes 'adaptive dynamics', which is a framework that has recently proved successful for modelling speciation events. I then explain why this framework cannot be successful as a solution for this particular problem.

Section 6.2 defines the intraspecific competition model and 6.3 describes some basic features of the model's behaviour. The important features are:

- (1) As long as both species survive early density fluctuations, they will usually stabilise with equal densities.
- (2) As long as the two species are reasonably competitive, speciation tends to occur in at least one of the two, and cyclic competition will occur among the resulting morphs.

The reasons for these two phenomena are explained further in chapters 7 and

8 respectively.

Section 6.4 assesses the model specifically in initial two-species states that could be considered plausible in nature. In these cases, the emergence of a cyclic ecosystem is unlikely under the model unless the intergenerational rate of phenotypic change is high.

Finally, section 6.5 explains why the model's assumptions are so restrictive that it cannot serve as a general model of intraspecific competition in natural ecosystems. This conclusion motivates a revision to the model which forms the basis of chapter 9.

6.1 Modelling speciation

The question of how RPS and other non-transitive cyclic competition systems become established has not yet been answered. Recent work has shown that it is possible for ecosystems containing more than three species to collapse, through one or more extinction events, into a RPS system (Cranefield, 2001). In this chapter I look at the possibility of a RPS ecosystem evolving, via a speciation event, from a simpler two-species ecosystem.

In three-species cycles, there is an asymmetry in the competitiveness of any two of those species in the absence of the third. If any of the three species is removed from RPS, its predator species goes extinct soon afterwards.

Therefore, any two-species system from which RPS can evolve must be unstable. Because of the instability of this pre-existing system, a third species must appear at just the right time in order to stabilise it, or else it will quickly collapse. This may be part of the reason why few RPS ecosystems have been found in nature.

In this chapter I examine a model which shows that the simple addition of an evolving intraspecific competitiveness trait to the pre-existing population results in exactly the kind of speciation event that stabilises the system.

6.1.1 Sympatric speciation

An increase in species diversity is usually thought to require either the physical isolation of two populations of a species (allopatric speciation), or the occurrence of an evolutionary branching event in a mixing population (sym-

6.1. MODELLING SPECIATION

patric speciation). Recent work in theoretical ecology has suggested that sympatric speciation may be responsible for most currently observed species diversity (Dieckmann and Doebeli, 1999; Day and Young, 2004; Doebeli and Dieckmann, 2004).

Before going on to discuss the new model in the next section, I will summarise the essential features of a class of mathematical models called adaptive dynamics (Dieckmann, 1997; Metz et al., 1996) that have recently been very successful in modelling sympatric speciation events. I will then explain why the framework of these models renders them unable to account for the origin of RPS.

6.1.2 Assumptions of adaptive dynamics models

An evolutionary branching event is a process in which a relatively homogeneous population splits into two distinct phenotypes (Doebeli and Dieckmann, 2000). Adaptive dynamics (AD) models attempt to model such a process using a very simple population structure.

Infinite population

Numbers of individuals of a species are typically described in terms of proportions of the total population, so the population size is effectively infinite as in the mean field model of RPS.

Single phenotypic trait

Individuals vary only in respect of a single phenotypic trait. In most AD models, this trait value is represented by a continuous variable even though the underlying genetics is discrete. The variable therefore represents a complex trait determined by the addition of small effects at a large number of genetic loci, each of which has two (or few) alleles, or a trait determined by a large number of possible alleles at a single locus.

Initially, populations are assumed to be monomorphic for the trait, meaning that all individuals are clustered tightly around a single mean phenotypic value.

Disruptive selection

Evolutionary branching points are found by looking for the points at which a population faces 'disruptive selection', or selective pressure in opposite directions on a phenotypic trait (Bell, 1997) at a single instant. In a monomorphic population facing disruptive selection, individuals with trait values on either side of the mean will be fitter than those whose trait values are equal to the mean. Other things being equal, the distribution of phenotypes will slowly become bimodal as the more extreme phenotypes reproduce more rapidly than the mean-valued phenotypes.*

Asexual population or assortative mating

Because the trait is additive, random pairings in sexual populations will tend to produce intermediate phenotypes, so unless there is assortative mating, mean-valued phenotypes will continually appear, extreme trait values will be diluted and the population will remain monomorphic (Seger, 1985; Dieckmann and Doebeli, 1999).[†]

While some AD models have been extended to model sexual populations (Kisdi and Geritz, 1999; Bürger, 2002a,b), the majority restrict their analysis to asexual populations, as does the SCA model presented in section 6.2. This is primarily for simplicity, but it also has the advantage that mutations which take the phenotype in one direction through trait space can accumulate over time, otherwise speciation would be impossible.[‡]

^{*} Disruptive selection is not just a theoretical construct: it has been observed in real populations. For example, among sticklebacks, a fish which inhabits Canadian lakes, disruptive selection is caused by intraspecific competition for two different types of resource. Limnetic sticklebacks live in the open water and eat plankton. They are morphologically different to benthic sticklebacks, which live in the vegetated parts of lake and eat bottom-dwelling prey. Under heavy intraspecific competition, monomorphic stickleback populations (in which intermediate individuals are able to use both resource types moderately well) have been shown to become dimorphic (with each morph becoming specialised in the use of one resource) through disruptive selection (Lavin and McPhail, 1985; Schluter and McPhail, 1992; Bolnick, 2004).

[†]In the case of the stickleback, for example, dimorphic populations are maintained through mate choice on behalf of both the male and female (Ridgeway and McPhail, 1984).

[‡]There is disagreement in the biological literature on exactly what counts as a species for asexually reproducing organisms (for example see Ghiselin (1987); Hull (1988)). For the purposes of the single-trait models of evolutionary branching discussed in this chapter, 'species' is used to refer to any group of individuals which is similar with respect to the phenotypic trait in question, and which is also related by descent.

6.1.3 Evolution in AD models

The properties described in section 6.1.2 are not particularly unique to the AD models; many of these assumptions will also be used in the SCA model. The defining assumption of AD models is the way in which they treat the evolutionary process.

AD models assume that mutations are rare, so a population always has time to reach an equilibrium state between the introduction of new mutations. From the initial state in which individuals of all species cluster around a single phenotypic trait value, the evolutionary process proceeds by repeating the following steps:

- (1) a rare mutant phenotype is introduced whose trait value is close to the mean, and then
- (2) the population densities are allowed to reach a stable state given the trait values of the mutant and any resident phenotypes.

Usually the result of step (2) is that the mean phenotype changes to that of the mutant, or the mutant dies out straight away.

The assumption that densities reach equilibrium between each mutation allows the direction of the selective force to be found by examining the way a mutant's fitness changes with respect to its trait value. The details of how the evolutionary branching points are found are described in the appendix, because the concept of invasion fitness used in that method turns out to be useful in the explanation (in chapter 7) of an aspect of the behaviour of the model in the next section.

6.1.4 Adaptive dynamics and RPS systems

AD methods are unsuitable for finding evolutionary branching points in an asymmetrical two-species competition system because AD makes the assumption that the population densities reach a stable point between mutations. In other words, it assumes a complete separation between the evolutionary timescale at which the phenotypic trait values change, and the ecological timescale at which the population densities change. The true situation will be approximated by such a separation only when mutations are very rare, or when ecological interactions and invasions are very frequent.

The emergence of RPS is a problem precisely because any two-species system

will be ecologically unstable up until the point at which one of the species goes extinct. If the AD timescale assumption were applied, an extinction of one species would occur in step (2), before any mutations could arise.

Therefore, emergence of RPS from two species requires that the ecological and evolutionary processes be able to work in parallel. This is a characteristic of the model introduced in the next section.

6.2 The intraspecific competition model

I describe a version of an unpublished cellular automaton model devised by Marcus Frean, in which an individual trait for intraspecific (within-species) competition is assumed to vary across the population, and to vary continuously.

The motivation for the model is to show how such a trait may diverge as extreme valued phenotypes are favoured by selection while intermediate valued phenotypes are selected against. These results are presented in the next section.

6.2.1 State description

The model uses an SCA grid with periodic boundaries, representing a spatial landscape. Each site on the grid is occupied by an individual, and each individual belongs to one of two species numbered 0 and 1.

Intrinsic species-determined competitive ability

Each individual has an intrinsic level of competitiveness which is completely determined by its species, either c_0 for species 0, or c_1 for species 1. Both of these species-level competitive abilities are assumed to be greater than zero, and can be thought of as the relative efficiency with which each species converts its available resources into competitive energy. Because this is the sort of property that is unable to evolve quickly, c_0 and c_1 remain constant throughout the period of a simulation.



Table 6.1: The outcome of a single interaction is determined by the probability $p_{i \rightarrow j}$ that an individual *i* invades a site occupied by another individual *j*. s_i and s_j refer to the species of individuals *i* and *j* respectively.

Individual variation in degree of within- or between-species competition

Every individual varies in its intraspecific competitiveness x, which is assumed to be an additive, polygenic trait ranging between zero and one. If an individual *i*'s intraspecific competitiveness $x_i = 0$, *i* can only compete against individuals from the other species, if $x_i = 1$, *i* can only compete against other individuals of its own species, and if $x_i = \frac{1}{2}$, *i* is equally competitive against all others regardless of their species. So each individual faces a very simple linear tradeoff between using its competitive resources for within- or between species competition.

6.2.2 Ecological interactions

The ecological process is performed by repeatedly choosing a grid site and one of its four (north, south, east, west) neighbours at random to interact. In a single interaction between two neighbouring individuals *i* and *j*, the chance $p_{i\rightarrow j}$ of *i* invading *j* is the competitive ability of *i*'s species multiplied by either *i*'s *intra*specific competitiveness x_i (if *i* and *j* are the same species), or by *i*'s *inter*specific competitiveness $1 - x_i$ (if *i* and *j* are different species). These probabilities are summarised in table 6.1(a). $p_{i\rightarrow j}$ is *i*'s chance of invading *j* once *i* has been chosen as a potential attacker and *j* as a potential victim, so this probability is always positive and nonzero even when *j* is more competitive than *i*.

When describing the behaviour of the system, a 'generation' of the simulation will refer to a number of individual interactions equal to the number of sites in the grid.

6.2.3 The evolutionary process

Upon a successful invasion, the victim's site becomes occupied by a new 'child' of the invader. The child's level of intraspecific competitiveness is calculated from the parent's by adding a small Gaussian mutation (with standard deviation σ), and then restricting the result so that it remains within the range [0,1].

6.2.4 Initial state description

The grid sites are assigned randomly at the beginning of the simulation. In general, the initial state of a simulation can be described by the following quantities:

<i>c</i> ₀ , <i>c</i> ₁ :	The competitive abilities of the two species, which are assumed not to change over time,
$ ho_0(0), ho_0(1)$:	The initial densities of the two species,
$p(x_0), p(x_1)$:	The initial distributions of the intraspecific competitiveness traits for individuals of each species, and
σ:	The standard deviation of the mutations.

Simplified initial state

For the rest of this section it is assumed that one species is capable of outcompeting (or competing equally against) the other, and I will call the stronger species *A* and the weaker species *B*. The parameter β describes the ratio of *B*'s competitive ability c_b to *A*'s competitive ability c_a . It will be assumed that $c_a = 1$, that $c_a \ge c_b > 0$, and therefore $0 < \beta \le 1$. The invasion probabilities $p_{i \rightarrow j}$ are shown in terms of β in table 6.1(b).

Also, because the entire lattice is always completely occupied, the two species densities must sum to one throughout the simulation, so the single quantity ρ will be used to refer to the density of species A, with ρ_0 being A's density at the start of the simulation. Throughout the rest of this chapter it will be convenient to drop the parentheses and refer to the variable densities of A and B during the simulation as ρ_A and ρ_B . Because these could be confused with the initial density ρ_0 , I will always use the term 'initial densities' to disambiguate.

As a result of these simplifications a simulation's initial state can now be characterised using two fixed quantities β and σ , the initial density ρ_0 , and the two initial trait distributions $p(x_A)$, $p(x_B)$.

6.3 Behaviour of the model

There are a number of features of the model which are not obvious from its definition. These are summarised in the rest of this section.

6.3.1 Stable coexistence

It is very common for both species to coexist indefinitely, even in cases when the relative strengths of the two species are unequal.

Whether or not one of the two species goes extinct depends on the initial conditions. Figure 6.1 shows the winning species on an 80 × 80 grid when the initial trait distributions are chosen randomly, and $\sigma = 0.01$. Coexistence is the most common outcome, and would be more likely on a larger grid.

Extinctions appear at the edges of the figure (shown by the yellow and light blue dots), either when species strengths are very uneven (β is small), or when the initial species densities are very uneven (ρ_0 is close to zero or one).

For by far the majority of this parameter space, if there is coexistence after 10,000 generations, then the species densities at equilibrium are close to half and half — indicating a very low risk of an extinction occurring. The robust stability of the coexistent equilibrium state at $\rho = 1/2$ is independent of whether interactions in the model are local or long-range, and this implies that it is a different sort of stability to that of the RPS model. The causes of the model's stability are analysed in more detail in chapter 7.

6.3.2 Early oscillations in species densities

It is not surprising that when one species is very much stronger than the other (β is very small), it is the weaker species that is likely to go extinct, which is why there is a yellow band at the bottom of figure 6.1.

But the winning species is unpredictable when the extinction is caused by unevenness in the initial densities, because the more vulnerable species is not always the one that is initially rare. This is shown by the combination of light blue dots (species B) and yellow dots (species A) at the left and right



Figure 6.1: The winning species or final density of species *A* after 10,000 generations on an 80 × 80 grid when $p(x_A)$, $p(x_B)$ are uniform distributions. Each point shows the result of one simulation for a particular combination of β and an initial density of *B*s. The yellow and light blue points indicate saturation by *A*s and *B*s respectively; all other colours, including black, indicate coexistence after 10,000 generations. In the regions of coexistence the shading from blue through black to orange represents the final species densities: black is close to half *A*, half *B*, blue is close to all *B* and orange is close to all *A*.


Figure 6.2: Large fluctuations in species densities at the beginning of a simulation. The graph shows the species densities ρ_A , ρ_B and the mean levels of intraspecific competitiveness for the first 250 generations of a 1000 × 1000 grid simulation for initial species densities (0.18,0.82), $\beta = 0.6$, $\sigma = 0.01$, and uniform initial distributions $p(x_A)$, $p(x_b)$.

edges of the figure. The behaviour is a result of oscillations in species densities which occur when the initial densities are uneven. An example of the density oscillations is shown in figure 6.2.

When species *A* is initially very common, species *A* individuals tend towards intraspecific competition very quickly because there is very little chance of meeting an individual of species *B*. Species *B*, which now faces very little competition from species *A*, increases in population size until it becomes the common species, which in turn leads to a tendency to intraspecific competition in species *B*, and so on. For both species, high levels of intraspecific competition (high *x*) follow periods of high species density (ρ) with a short time lag. Similarly, periods of high interspecific competition (low *x*) follow periods of low species densities.

Vertical bands of extinctions of the same colour are visible on the left and right hand sides of figure 6.1, and these are due to the timing of the oscillations. This can be confirmed by noting that there is an order in which the species reach their minimum densities. Even though species *A* is initially rarest in figure 6.2, the initial trajectories of the densities ensure that species *B* is in fact more likely to go extinct, because it reaches a lower minimum density. Its lowest density of 0.10 occurs after 26 generations at the point marked min(ρ_B), and *A*'s lowest density of 0.16 occurs after 95 generations at the point marked $\min(\rho_A).$

6.3.3 Evolutionary branching

In many cases an evolutionary branching event occurs in one or both species, in which the values of the traits of the individuals diverge. This produces a competitive cyclic ecosystem with three or four distinct morphs. Individual competitiveness tends towards extreme interspecific or extreme intraspecific competition; distributions of x bunch up at zero or one and individuals with intermediate values of x die out.

Whether the branching event occurs depends primarily on the initial state of the simulation; there is very little variation due to randomness. Several other possible outcomes are described in section 6.3.4.

An example of the emergence of RPS cycles is shown in figure 6.3(a), where a simulation of 250,000 individuals begins with uniform distributions of x, and ends with extreme values of x for most individuals.

The final distribution of x is stable and an attractor. For example, even if the simulation begins with fixed intermediate values for all the x_i , as shown in figure 6.3(b), the x_i values still tend towards zero and one and the two simulations have converged to similar patterns after 2500 generations. The equilibrium state has the same dynamics as the three-species spatial RPS game, with intransitive competition occurring between interspecific A, interspecific B, and intraspecific A. Interspecific A beats interspecific B, which beats intraspecific A, which in turn beats interspecific A.

Figure 6.4 shows the trait distributions over time. Initially, individuals of both

Figure 6.3: SCA simulations of the model described in section 6.2. The pictures show the composition of a 500 × 500 grid over time with $\beta = 0.6$, $\rho_i = 0.5$, $\sigma = 0.01$ for two different initial distributions of the individual levels of intraspecific competition x_i . Sites occupied by the stronger species A are shaded from dark brown for interspecific competitors (x = 0) through to orange for intraspecific competitors (x = 1), and sites occupied by the weaker species B are shaded similarly from dark to light blue. In (a), the initial x_i s are uniformly distributed in [0,1], and in (b), initial x_i s are fixed at 0.4 for species A, and at 0 for species B. In both cases, intermediate shades of brown and blue disappear and eventually only separate regions of orange, dark brown and dark blue remain, indicating a tendency in species A to move away from a mixture of interspecific competition and intraspecific competition to a focus on one or the other exclusively, and a tendency in species B towards interspecific competition.







Figure 6.4: Histograms showing the distribution of inter- and intraspecific competitors of each species over time for a simulation on a 1000 × 1000 grid with $\beta = 0.6$, $\rho_i = 0.5$, $\sigma = 0.01$ and a uniform initial distribution of *x*.

species are uniformly distributed, but slowly bunch up at zero and one. In this example, the distributions at time t = 3000 are stable. Some individuals with x between 0 and 1, are always present in the final distribution, because ongoing invasion and mutation means there will always be individuals who have mutated away from the optimum. The number of such individuals present in the final distribution depends on the size of the mutations σ (see figure 6.5).

While the stability of the coexistent state occurs regardless of whether interactions are local or long-range, evolutionary branching in the population is very



Figure 6.5: Distributions of *x* for each species after 15,000 generations on a 1000 \times 1000 grid for several values of β and three values of the mutation size σ .



Figure 6.6: Net invasion rates among the pure strategies where x = 0 or x = 1. A0 and B0 are interspecific (x = 0) competitors of species A and B respectively, and A1 and B1 are intraspecific (x = 1) competitors of species A and B respectively. The labels on the arrows are the net invasion rates between these strategies from table 6.1. (a) shows the net invasion rates between the four strategies that are present when β is high, and (b) shows the net invasion rates between the three strategies that are present when β is at an intermediate level.

much dependent on local dispersal. Dimorphic populations like the species *A* population in figure 6.4 only occur when there is some clustering of similar individuals on the landscape. I suggest some reasons for this dependency in chapter 8.

6.3.4 Attracting states in the model

The system frequently reaches a stable state, but not always the RPS-like state of figure 6.3. Sometimes there is a branching event in both species, and sometimes both species remain monomorphic.

With enough simulations, it is possible to see what kinds of attractors the system falls into at equilibrium, as a function of the model parameters. The simulations are initialised with equal densities of each species and uniform trait distributions, and run for a long time, after which the distribution of inter- and intraspecific competitors of each species is plotted. This results in a set of distributions like those shown in figure 6.5.

When the two species are evenly matched (β close to 1), there is a branching event in both species. If β is below a threshold around 0.85, and above about 0.2, only the stronger species splits into two morphs while the weaker one stays interspecific only. There are a range of equilibrium distributions occu-



Figure 6.7: Variation in equilibrium trait distributions with β and σ . A total of 1000 simulations were run on 1500 × 1500 grids, for ten values of σ and 100 values of β , with trait distributions initially uniform and with equal initial species densities. The trait distributions after 15,000 generations were then classified into seven types (a)-(g) described in the text.

pying the region where the species are very mismatched (β less than about 0.2), but if β is reduced far enough, the result is a two-morph system in which the strong species is fully intraspecific and the weak species fully interspecific.

The reasons for this can be partially explained by considering the net invasion rates among the 'pure-strategy' morphs in which x = 0 or x = 1, shown in figure 6.6. Figure 6.6(a) shows the situation among four morphs. As β is reduced, *B*1 invades *B*0 more slowly and *A*0 invades *B*0 more quickly, cutting *B*1 out of the loop by predating on its resource. *B*1 is thus the most vulnerable of the four pure strategies. Among the three remaining pure strategies shown in figure 6.6(b), the situation is like RPS in which we have seen in section 1.1 that the density of a species is proportional to the invasion rate of its prey. As β is reduced further, *A*0 becomes the most vulnerable because its prey *B*0 has





Figure 6.8: Examples of the state of the grid for the main categories (a)-(f) of attracting equilibrium distribution described in the text. The arrows below the histograms show the net frequency of changes in occupant over time for ten randomly selected cells, where the trait space has been divided into ten equal-sized bins in [0,1]. Darker arrows show the most common transitions and lighter arrows the less common ones. Transitions between bins that are very close to one another have been removed.

the smallest invasion rate.

Mutation size

The effect of the size of mutations, σ , is visible in the figure 6.5 distributions. When σ is small, child phenotypes remain close to the parent and individual phenotypes are tightly bunched up around the mean of the morph. When σ is large the mutations move children farther away, so the distributions are more spread out.

Distribution categories

Figure 6.7 shows a classification of the equilibrium distributions into seven categories (labelled (a)-(g) in the figure) depending on whether each species

had concentrations of individuals with *x* close to, zero, one, or at some intermediate value.

The characteristics of the categories are described below, and examples of each category are shown in figure 6.8.

- **Case (a):** Figure 6.8(a). When β is high, the equilibrium distribution is composed of four main clusters of individuals (highly intraspecific and highly interspecific competitors of both species) in a nontransitive, cyclic competitive configuration among four groups. Interspecific *A* invades intraspecific *B* invades interspecific *B* invades intraspecific *A*.
- **Case (b):** Figure 6.8(b). The 'three-species' RPS dynamics of figure 6.3. The stronger species *A* splits into extreme interspecific and extreme intraspecific morphs, while species *B* concentrates on interspecific competition.
- **Case (c):** Figure 6.8(c). Species *A* has one concentration at x = 1 and another intermediate concentration, while species *B* concentrates on interspecific competition, sometimes with a small number of intraspecific competitors as well.
- **Case (d):** Figure 6.8(d). Species *A* concentrates on intraspecific competition while species *B* has two established concentrations of individuals with x = 0 and x = 1.
- **Case (e):** Figure 6.8(e). Species *A* concentrates on intraspecific competition while species *B* has a concentration of individuals with intermediate *x*, and also sometimes smaller concentrations near x = 0 and x = 1.
- **Case (f):** Figure 6.8(f). The population consists primarily of two groups, interspecific *B* and intraspecific *A*, and a relatively stable spatial distribution with a very low turnover of species in most regions of the grid.
- **Case (g):** The weaker species goes extinct when β is close to zero, resulting in a monoculture of intraspecific *A*.

In very large portions of the parameter space, cases (a)-(d) in figure 6.7, the system reaches a state with at least one dimorphic population and intransitive competition.

6.4 Realistic scenarios

In the last section, most simulations started with equal species densities and large variations in intraspecific competitiveness. Those initial states made it easier to assess the equilibrium model behaviour, but they were also unrealistic. If more plausible initial population structures are considered, early extinctions become much more likely.

Unequal initial species densities

When two competing species initially come into contact, it is unreasonable to assume that they could be equal in number or randomly-mixed. It is more likely that the contact begins with a single immigrant, or a very small group of immigrants invading a large resident population, so I will consider these scenarios only.

There is little point in examining the case when a new immigrant is at a disadvantage and goes extinct quickly. I consider the situation where single immigrant *i* has a higher invasion probability versus the resident *j* than the resident has versus the immigrant, $p_{i\rightarrow j} > p_{j\rightarrow i}$.

Extreme intraspecific competitiveness levels

Populations come together out of geographic isolation only rarely. So it is likely that both the population from which the new immigrant is drawn, and the resident population, will have reached an equilibrium state at the time of the introduction. Therefore I consider populations that are either fully interspecific or fully intraspecific, because that is the equilibrium state of the majority of populations that are subject to local spatial effects.

6.4.1 Scenario types

There are three possible scenarios for the meeting of two monomorphic, extreme populations:

(1) Two relatively evenly matched species are competing interspecifically, but one is slightly stronger and is slowly driving the other towards extinction. The immigrant is a single individual of the stronger species *A*, the resident population is made up of *B*s. Everyone's intraspecific competitiveness *x_i* starts at zero, so the immigrant invades the resident at a

net rate of $1 - \beta$. Stable coexistence requires the immigrant to expand at the start and to evolve intraspecific competitiveness before the resident is wiped out.

- (2) A resident population, which has become extremely intraspecific in the absence of a competitor, is being invaded by a more interspecific immigrant population. The immigrant is an interspecific species *B* individual, and the resident population is made up of intraspecific species *A* individuals, so the immigrant invades the resident with a net rate of β . Stable coexistence requires the resident to evolve interspecific competition before it is wiped out by the invading immigrant.
- (3) Two isolated intraspecific populations come into contact, but they cannot compete with one another directly until one evolves interspecific competitiveness. In this case the species do not initially compete, so there is plenty of time for an interspecific mutant to arise.

While scenario 3 may be plausible, it turns out to be similar to scenario 2. There is plenty of time for the initial interspecific mutant to appear in one of the two species. But after that the system will be unstable, because the new interspecific mutant of one species will wipe out the second species entirely unless an interspecific mutant of the second species appears quite quickly. While waiting for the second mutant, the situation is similar to scenario 2, where an interspecific immigrant is driving out an intraspecific resident. I will therefore ignore scenario 3 and present results of simulations of scenarios 1 and 2.

Of these remaining scenarios, I consider the second to be the most plausible, because an isolated species with no competitors is likely to be fully intraspecific. I will also look at the behaviour of the first, because it possible for species to be 'pre-adapted' to a potential competitor for other reasons.

6.4.2 Simulations of scenarios 1 and 2

In scenarios 1 and 2 one of the two species is always represented by a single individual, so the initial densities are very uneven. We have already seen in figure 6.1 (page 96) that extinctions are common but not inevitable in such conditions.

Figure 6.9 shows the densities and mean trait values for examples of simulations which don't result in extinctions. The scenario 1 example begins farther from equilibrium, and there is a massive spike in the resident's intraspecific



Figure 6.9: Examples of simulations on a 1000×1000 grid leading to stable coexistence from the starting conditions of scenarios 1 and 2 in section 6.4.1. The solid lines show the densities of the two species, and the dotted lines their mean levels of intraspecific competitiveness. The scenario 1 simulations start with one species *A* individual with x = 0.01, and 999,999 species *B* individuals with x = 0.01. The scenario 2 simulations start with one species *A* individuals with x = 0.99.

competitiveness before it settles back down to its equilibrium level once the species densities even out. In the scenario 2 example, the adjustment happens smoothly. In both cases, the densities are slow to change, so the system spends a long time in a state where the immigrant species is at risk because of its low numbers.

Effect of β and σ

Whether the system stabilises depends on both β and σ . The total population size has an effect too, but it is less important once the population is sufficiently large.

The invasion speed of the immigrant is determined by β . In scenario 1, invasion is slowest when β is high, and in scenario 2, invasion is slowest when β is low. Slow invasion does not always increase the likelihood of coexistence, however, because a slow invader spends more time at risk in the early stage with very low numbers.

The responsiveness of the resident to the invader is determined by σ . In both scenarios, high σ gives a greater likelihood of coexistence.

This dependence on β and σ is pictured in figure 6.10 which shows results of simulations of both scenarios for 22,500 different combinations of β and σ .

Extinctions in scenario 1

In scenario 1, species *A* is initially rare, but is stronger and invades species *B* at a net rate of $1 - \beta$. Some simulations are won by species *B* just because species *A* is initially so rare that it is wiped out at the beginning through bad luck, before it has time to expand. These are the earliest extinctions, and they appear as scattered dark blue dots in the top picture in figure 6.10. There are more of these extinctions when β is high, because *A*'s population expands faster when β is low.

Secondly, some simulations are won by species *A* because *A* has a higher invasion rate, and wipes *B* out before *A*'s intraspecific competitiveness levels are able to respond to *A*'s increasing population. These simulations appear as a brown triangular region at the bottom left of the top picture. σ has an important effect on these secondary extinctions; whether species *A* wins appears to depend on a linear combination of β and σ .

There is also a dark brown strip at the bottom of the diagram. This is the region labelled (g) in figure 6.7. In this region *B always* goes extinct at equilibrium, even when the initial densities are even and the initial x_i are uniformly distributed. The orange area at its border is region (f) in which *A* remains very intraspecifically competitive, *B* very interspecifically competitive but the grid is quite static.

110





Figure 6.10: Regions of extinctions and coexistence in the β - σ parameter space for simulations starting with the conditions of scenarios 1 and 2. Each point represents the state of a SCA simulation on a 300 × 300 grid after 10,000 generations. If the simulation ended with the extinction of one species, the winning species is marked by a either dark blue point (for species *B*) or a brown point (for species *A*). All other colours shown on the colour scale, from light blue through to white and then orange indicate the species densities after 10,000 generations.

Extinctions in scenario 2

In scenario 2, species *B* is initially rare and invades species *A* at a net rate of β . The earliest extinctions happen when species *A* wipes out species *B* before it has the chance to expand, and appear in the bottom diagram of figure 6.10 as brown dots scattered over the whole picture. They are more common when β is low because that is when species *B* expands slowest.

The early extinctions are much less common than the ones in scenario 1 because in this case, species *A* individuals start out as almost fully intraspecific, so they have very little chance of winning any interactions with species *B* individuals. Even when the *net* invasion rate β of species *B* in scenario 2 is higher than the net invasion rate $1 - \beta$ of species *A* in scenario 1, these early extinctions are less common in scenario 2 because of the difference in the initial *absolute* invasion probabilities.

The second wave of extinctions occurs when species *B* slowly takes over the grid before the intraspecific competitiveness levels have time to respond and regulate the system, and these simulations appear as a blue triangle in the upper left of the picture. These secondary extinctions are strongly affected by σ as well as β .

The dark brown region (g) is the same shape as region (g) for scenario 1, which is not surprising because these represent simulations that have gone past the early extinction period to an equilibrium that does not depend on the initial state of the grid.

6.4.3 Plausibility of mutation rates

While there are few limits on the values that β could take in real ecosystems, σ is different because inter-generational phenotypic change is measurable. It might be objected that the values on the *x*-axis in figure 6.10 are simply implausible as mutation rates, given that genetic mutations have been estimated to occur on the order of 10^{-9} per base pair per year in mammalian genomes (Kumar and Subramanian, 2002), and 10^{-3} per genome per replication (or cell-division) in all DNA-based genomes (Drake et al., 1998).

However, there are many complicating factors. For example, mutations are known to occur at very different rates in different genes (Wolfe et al., 1989). There are also known mechanisms of both prokaryotes and simpler eukaryotes which allow faster evolution when populations are placed under stress, for example when confronted with a host's immune response, or starvation (Moxon et al., 1994; Adams et al., 1998). Some of these mechanisms can increase the bacterial mutation rate by a factor of 10^4 (Miller, 1996; Bjedov et al., 2003).

Even though I have often referred to σ as the 'mutation rate' or the 'size of mutations', it does not represent genetic mutations alone. It includes all the ways in which a phenotype can change from one generation to the next: such as the genetic mutation in all the alleles that make up the trait, crossover (for sexual species), and potentially even environmental elements such as learning which can be passed on to children. A large number of mechanisms are involved (Poole et al., 2003), and σ is the sum of all these and represents the total 'evolvability' of the two species.

6.5 The intraspecific competitiveness trait

An important objection to the intraspecific competition model is the lack of realism in the way the trait affects individual interactions. In section 6.2, x_i is described as "*i*'s level of intraspecific competitiveness". However, there is no general form that intraspecific competitiveness takes, and in the real world such a trait will differ greatly from species to species. In this section I will explain why although many biological traits play a part in intraspecific competition, few of these traits are likely to fit the model's assumptions used in determining the invasion probabilities.

Assumptions of the model

The crucial assumptions about the invasion probabilities are contained in table 6.1 on page 93, and they place quite specific restrictions on the way the *x* trait affects individual competitive interactions. The definition of $p_{i\rightarrow j}$ implies that

- (1) Some aspects of competitiveness are determined by the *x* trait, for which individuals face a smooth, linear tradeoff; but
- (2) Other aspects of competitiveness are determined by the β ratio and these aspects depend on the individual's species, and don't change perceptibly during the timeframe under consideration.
- (3) An individual's ability to invade another's territory is largely determined by its trait, but the same individual's susceptibility to being colo-

nised is determined by its species.

(1) Tradeoff between intraspecific and interspecific competition

The tradeoff is inherent in the fact that *i*'s intraspecific competitiveness is defined as x_i and *i*'s interspecific competitiveness by $1 - x_i$. On its own, this assumption is reasonable; there are lots of traits under the control of a large number of genes, which will be subject to evolutionary tradeoffs. An example is the tradeoff facing sticklebacks (see the footnote on page 90) over adaptation to a particular niche. Competitiveness versus particular species is bound to be affected by such traits.

(2) Species-determined competitiveness

The asymmetric competitiveness of the two species, represented by the ratio β , is more restrictive but not unrealistic. Competitive asymmetry between species is common in the natural world, and the competition is not always in direct contests — some species are just better adapted at using resources efficiently than others.

In the model, it is assumed that this species-determined competitive ability is independent of any individual's intraspecific competitiveness trait, so some very fundamental characteristics such as the species' basic morphology or resource niche must be relatively immutable during the time periods under consideration. If a species' primary nutrients are difficult to change, while its method of territorial colonisation is more plastic, then such a species is represented well by the second assumption above.

(3) An individual's susceptibility to specialists is not dependent on its own specialisation

The third assumption is problematic. This is the assumption that all the characteristics that go to make up individual i's susceptibility to colonisation are based wholly on i's species, and hence must be independent of i's investment in intraspecific or interspecific competitiveness.

This is manifested in the model in table 6.1(b). The individually-varying factor of an individual's invasion probability (whether x_i or $1 - x_i$ appears in the third column) depends on *j*'s species, while the species-dependent factor (whether β appears in the third column) depends on *i*'s species.

This restriction probably rules out most morphological characters of animals

and plants because it is very difficult to imagine a character which has a large effect on 'attack', but negligible effect on 'defence'. For example, body size in animal contests is not the sort of trait that can be modelled by the system. Even if large body size is advantageous in contests against species X and small body size is advantageous against species Y, for the model to be accurate, having a large body size would have to be equally advantageous against both large and small individuals of species X, and small body size would have to be equally advantageous against both large and small individuals of species X.

6.6 Summary

The adaptive dynamics framework cannot be used to explain the evolution of RPS ecosystems, because it assumes a complete separation of the evolutionary and ecological timescales.

A two-species CA model with individual variation in intraspecific competitiveness can become a RPS ecosystem after a speciation event which occurs in a large portion of the parameter space.

However, when more realistic two-species ecosystems are used as starting points, early extinctions are very common. The emergence of cyclic competitive ecosystems will require high rates of per-generation phenotypic change, or species whose intrinsic levels of competitiveness are not especially uneven.

Also, the model can only apply when individuals' method of attack or colonisation is species-independent, while at the same time individuals' susceptibility to attack is species-dependent. This assumption limits the model's potential applicability to a narrow range of potential traits, and consequently the model fails to represent intraspecific competition in general.

Toxin-producing organisms are a more promising area of application, because toxins are potentially effective against specific species while at the same time leaving others unaffected: susceptibility to toxins is often species-dependent. An adaptation of the model for ecosystems involving toxin-producers is explored in chapter 9. CHAPTER 6. INTRASPECIFIC COMPETITION

Chapter 7

Stability of species densities in the intraspecific competition model

One of the aspects of the intraspecific competition model of chapter 6 that requires explanation is the near-universal stability of the two species densities around $\rho = \frac{1}{2}$ illustrated in figure 6.1 (page 96). The spatial RPS game is also very stable with three species (see section 1.2, page 4) so it might not seem surprising that the current model has stable species densities too, because we know that the intraspecific model acts in many ways like RPS.

But the stability arises from different causes in the two models. Community stability in the spatial RPS game is dependent on a spatially-structured population. When individuals play the game in a big pool (section 1.3, page 6), there is no state with stable densities in which more than one species survives.

In contrast, densities in the intraspecific competition model are usually stable in a well-mixed population (figure 7.1). Without spatial structure, there is stability at $\rho = \frac{1}{2}$, and a uniform distribution (when $\beta = 1$) of phenotypic trait values. The figure shows that although spatial structure is responsible for divergence in the individual character traits, it is not responsible for the stability in species densities.

In the remainder of this chapter I describe a simplified nonspatial, infinitepopulation model of the intraspecific competition system. By analysing this model, it is possible to show why the species densities are always centred around the value $\rho = \frac{1}{2}$.

Nonspatial, finite-population simulations are then used to argue for the claim that over-regulation by the traits attracts the densities to $\rho = \frac{1}{2}$.



Figure 7.1: Average over 20 simulation runs of the distribution of inter vs. intraspecific competitiveness after 15,000 generations, for two species, with no spatial interactions (top row, population size 2,250,000), and with spatial structure and local interactions (bottom row, 1500 × 1500 grid), for $\beta = 1$ and $\beta = 0.6$. In all four cases $\rho_0 = 0.5$, $\sigma = 0.01$, and the initial distributions $p(x_A)$ and $p(x_B)$ are normal, centred around 0.6 and 0.52 respectively with standard deviation 0.05.

Finally, I show that despite the lack of trait divergence in the nonspatial system, the effect of spatial structure on extinctions in the model is relatively minor.

7.1 A nonspatial intraspecific competition model

In this section I describe a numerical model of mean phenotypic values similar to the adaptive dynamics model, with a relaxation of the requirement for complete separation of the ecological and evolutionary timescales. With certain assumptions, it can be shown that the only factor affecting the direction of the trait values is whether the species density ρ is greater than $\frac{1}{2}$, or less than $\frac{1}{2}$.

7.1.1 Model definition

The population is made up of individuals from two species *A* and *B*, where *A* is stronger than *B* and the ratio of strengths is described by β as it is in the SCA model. Each individual *i* has a phenotypic trait z_i which determines, but is not equal to, its level of intraspecific competitiveness. To make the model easy to work with, I allow *z* to vary outside the range zero to one, but ensure that competitiveness levels stay between zero and one by scaling



Figure 7.2: The phenotypic effect of the linear trait value x of the computer model (dotted line) and the sigmoid trait value f(z) of the nonspatial model (solid line).

z with a sigmoid function. The function $f(z) = 1/(1 + \exp(2 - 4z))$ is used so that f(z) plays a similar role to the variable *x* in the computer model of section 6.2. The difference in the effects of *x* and f(z) are shown in figure 7.2, which shows that the two models concepts of intraspecific competitiveness are similar when *x* and *z* are close to $\frac{1}{2}$.

Individual *i*'s relative competitiveness versus *j* is defined as

$$\begin{aligned} f(z_i) - f(z_j) & \text{when } species(i) = species(j) = A, \\ (1 - f(z_i)) - \beta(1 - f(z_j)) & \text{when } species(i) = A, species(j) = B, \\ \beta(1 - f(z_i)) - (1 - f(z_j)) & \text{when } species(i) = B, species(j) = A, \\ \beta f(z_i) - \beta f(z_j) & \text{when } species(i) = species(j) = B, \end{aligned}$$

which are very similar to the invasion probabilities given in table 6.1 (page 93). The difference is that here the net invasion probabilities $(p_{i\rightarrow j} - p_{j\rightarrow i})$ are being described rather than the absolute ones, and $f(z_i)$, $f(z_j)$ have been substituted in place of x_i , x_j .

7.1.2 Ecological dynamics

First consider two mean phenotypes, \bar{z}_A and \bar{z}_B of species A and B respectively. Because the population is not spatially structured, the chance of meeting an individual of species A in a single interaction is ρ_A , and the chance of meeting an individual of species B is ρ_B . The following equations determine how the mean phenotypes perform against mean phenotypes:

$$\frac{d\rho_A}{dt} = \rho_A \rho_B \left(1 - f(\bar{z}_A) - \beta(1 - f(\bar{z}_B)) \right)$$
$$\frac{d\rho_B}{dt} = \rho_B \rho_A \left(\beta(1 - f(\bar{z}_B)) - 1 + f(\bar{z}_A) \right)$$

and because $\rho = \rho_A = 1 - \rho_B$, these can be simplified to

$$\frac{d\rho}{dt} = \rho(1-\rho) \left(1 - f(\bar{z}_A) - \beta(1-f(\bar{z}_B)) \right).$$
(7.1)

This equation simply describes how the densities of the two species would evolve in the absence of any mutations in \bar{z}_A , \bar{z}_B . There are uninteresting fixed points at $\rho = 0$ and $\rho = 1$, where one of the two species has gone extinct.

The densities are also fixed along the line at which $1 - f(\bar{z}_A) = \beta(1 - f(\bar{z}_B))$, where the mean levels of intraspecific competition exactly balance out the asymmetric growth rates of the two species. But this line is unstable in the sense that any slight deviation of \bar{z}_A or \bar{z}_B will push ρ towards one or zero.

7.1.3 Evolutionary dynamics

In the real world, the mean phenotypes will evolve through natural selection. Because selection can only happen through mutations, the adaptive dynamics method can be used, by examining the invasion fitnesses, or initial per capita growth rates r'_A and r'_B , of mutants z'_A and z'_B (see equation (A.2) on page 221). These growth rates depend on the difference between the mutant's level of intraspecific competitiveness and the mean levels of intraspecific competitiveness of each species in the population:

$$r'_{A} = \rho_{A} (f(z'_{A}) - f(\bar{z}_{A})) + \rho_{B} (1 - f(z'_{A}) - \beta(1 - f(\bar{z}_{B})))$$

$$r'_{B} = \rho_{A} (\beta(1 - f(z'_{B})) - 1 + f(\bar{z}_{A})) + \rho_{B} (\beta f(z'_{B}) - \beta f(\bar{z}_{B})).$$
(7.2)

So in the case of a species A mutant z'_A , the reasoning behind the equation for r'_A goes like this: if it meets another species A individual (which it does with chance ρ_A), its fitness is determined by how *intra*specifically competitive it is relative to its average opponent: $f(z'_A) - f(\bar{z}_A)$. If it meets a species B individual (chance ρ_B), its fitness is determined by how *inter*specifically competitive the mutant is: $1 - f(z'_A)$, compared to its average opponent: $1 - f(\bar{z}_B)$, moderated by their relative strength β . Because I assume the mutant is rare, there is no term to describe what happens if the mutant meets a copy of itself. A similar explanation applies for the species B mutant's invasion fitness r'_B .

The changes in these initial per capita growth rates with respect to the mutants' phenotypic trait values are

$$\frac{\partial r'_A}{\partial z'_A} = (2\rho - 1)f'(z'_A)$$

$$\frac{\partial r'_B}{\partial z'_B} = \beta(1 - 2\rho)f'(z'_B),$$
(7.3)

where

$$f'(x) = \frac{d}{dx}f(x) = \frac{4 \cdot \exp(2 - 4x)}{(1 + \exp(2 - 4x))^2}.$$

When equations (7.3) are evaluated at the mean phenotypes, $(z'_A = \bar{z}_A, z'_B = \bar{z}_B)$, the effects of selection on the mean phenotypes are obtained.

7.1.4 Evolutionary and ecological timescales

I assume that evolution drives the mean phenotypes proportionally to the gradients of the mutant growth rates with respect to mutant trait values in (7.3), and also assume that a ratio μ , describing the frequency of mutational change in the mean phenotypes to the frequency of ecological interactions, is constant.

A complete system can then be described by taking (7.1) and (7.3) together, and using the ratio μ to describe the relative difference between the two timescales.

The resulting model is defined by the following three equations:

$$\frac{d\rho}{dt} = \rho(1-\rho)\left(1-f(\bar{z}_A)-\beta(1-f(\bar{z}_B))\right)$$

$$\frac{d\bar{z}_A}{dt} = \mu \cdot \frac{\partial r'_A}{\partial z'_A}|_{z'_A = \bar{z}_A} = \mu(2\rho-1)f'(\bar{z}_A)$$

$$\frac{d\bar{z}_B}{dt} = \mu \cdot \frac{\partial r'_B}{\partial z'_B}|_{z'_B = \bar{z}_B} = \mu\beta(1-2\rho)f'(\bar{z}_B)$$
(7.4)

7.1.5 Summary of model assumptions

It is possible that the class of model described here will be useful in the general situation where the AD assumptions are too restrictive for the underlying population.

The model described in this section makes the following assumptions:

- (1) That the population is spatially unstructured, like the mean-field;
- (2) The population is infinite and only species densities are of interest;
- (3) Only the mean trait values are of interest;
- (4) Mutants are rare;
- (5) The evolutionary and ecological timescales operate in parallel with a constant relative speed.



Figure 7.3: The system defined by equations (7.4), for $\beta = 0.6$, $\mu = 0.01$. (a) Trajectories of ρ vs. $f(\bar{z}_A) - f(\bar{z}_B)$ for six initial values of $(\rho, \bar{z}_A, \bar{z}_B)$ in {(0.1,0.6,0.57), (0.18,0.6,0.52), (0.26,0.6,0.47), (0.34,0.6,0.42), (0.42,0.6,0.37), (0.5,0.6,0.32)}, with these initial values marked by red dots. The trajectory marked with a solid line in (a) appears in (b), plotted against time, along with the values of \bar{z}_A, \bar{z}_A .

7.2 **Properties of the nonspatial model**

In the system of equations (7.4) there is a fixed point when all three are equal to zero, which happens when $\rho = \frac{1}{2}$ and $1 - f(\bar{z}_A) = \beta(1 - f(\bar{z}_B))$. Around this point the system oscillates in a stable orbit centred on the fixed point, as shown in figure 7.3 — the fluctuations in density and mean trait levels continue with the same amplitudes forever.

Importance of $\rho = \frac{1}{2}$

The densities fluctuate around $\frac{1}{2}$, and the difference in mean levels of intraspecific competition fluctuate around the point at which $1 - f(\bar{z}_A) = \beta(1 - f(\bar{z}_B))$. The sign of the changes in trait values, whether they increase or decrease, depends only on whether $\rho > \frac{1}{2}$. This is because the only terms which



Figure 7.4: Behaviour of the system of equations (7.4) for variations in μ , the mutation rate of the mean trait values when $\beta = 0.6$. The red and blue lines indicate the mean trait values \bar{z}_A , and \bar{z}_B respectively, and the black line is the density of species $A(\rho)$. In all cases the initial values are $\rho = 0.35$, $\bar{z}_A = 0.7$, $\bar{z}_B = 0.6$.

can potentially go negative in the equations for $d\bar{z}_A/dt$ and $d\bar{z}_B/dt$ are the terms $2\rho - 1$ and $1 - 2\rho$ respectively.

So the densities are regulated by the intraspecific competition levels, with the periods of the oscillations in densities keeping time with the oscillations in intraspecific competition, for the same reasons as those described in section 6.3.2 on page 95 for the spatial simulation.

Importance of *µ*

The value of μ determines how quickly the intraspecific competitiveness trait is able to respond to changes in density. This is shown in figure 7.4. When $\mu = 0$, the slower species goes extinct because the faster one never becomes intraspecifically competitive. As μ increases, the oscillations become more frequent, the large amplitudes in the fluctuations of the densities become smaller, and the amplitudes of the fluctuations in the intraspecific competition levels increase.

7.3 Nonspatial finite-population simulations

In the nonspatial system of equations (7.4), the species densities do not converge to a stable point, instead they cycle endlessly. This is not what happens in a nonspatial simulation where the population is finite rather than infinite. Figure 7.5 shows the trajectory of ρ , \bar{x}_A , and \bar{x}_B for a nonspatial agent-based computer simulation. The simulation is of the intraspecific competition model described in section 6.2 (page 92), except that for each interaction two individuals are chosen at random from the population and need not be neighbours.

The trajectories shown in figures 7.5(a) and (b) correspond to the trajectory marked with the solid line in figure 7.3(a) and (b), which is the numerical result of equations (7.4) for the same parameter values. Comparing these figures shows that the numerical model is accurate at describing the initial behaviour in the agent-based simulations. However, in figures 7.5(c) and (d), where the same simulations are shown for an extended time period, the behaviour changes: eventually the agent-based model drifts toward stable values of ρ , \bar{x}_A , and \bar{x}_B , whereas in the numerical model the equivalent variables cycle forever in stable orbits.

In both nonspatial models, intraspecific competitiveness traits respond immediately to the densities crossing the point at $\rho = \frac{1}{2}$, (figures 7.3(b) and 7.5(b)) whereas in the spatial case there is a delay from the time at which species densities cross $\rho = \frac{1}{2}$ until the mean trait values change direction (figure 6.2, page 97). In the spatial model the global species densities cannot affect the majority of individuals straight away because individuals only experience population densities at the local level.

7.3.1 Intraspecific competition vs. RPS

The fact that densities are attracted to $\frac{1}{2}$ in the nonspatial finite-population system is not what one would expect based on the difference between the infinite and finite-population nonspatial models of the RPS system outlined in sections 1.1 and 1.3.

Infinite population models of both systems predict that densities cycle forever in stable orbits. But finite population models of the two systems show



Figure 7.5: Agent-based simulation without spatial interactions for a population size of 1,000,000. This is the agent-based simulation analog of the trajectory shown with the solid line in figure 7.3: $\beta = 0.6$, initial $\rho = 0.18$, and the initial x_i s are normally distributed (with standard deviation 0.05) around the means $\bar{x}_A = 0.6$, $\bar{x}_B = 0.52$. The size of mutations $\sigma = 0.01$. (a) shows the change in ρ vs. $\bar{x}_A - \bar{x}_B$, and (b) the change in ρ , \bar{x}_A , and \bar{x}_B vs. time for the first 250 generations, while (c) and (d) plot the same quantities for 15,000 generations. Figures (e) and (f) show a running average of \bar{x}_A and \bar{x}_B respectively (thick lines), taken over 150 generations for the first 8000 generations, with the thin lines showing a running average of one standard deviation from the means.

very different behaviour: in RPS, the species densities are unstable and grow increasingly different over time until one species becomes extinct. In the intraspecific competition model, the two species densities are attracted to $\frac{1}{2}$, and stable coexistence of both species is a likely outcome. So changing from infinite to a finite population has the opposite effect in the two ecosystems.

Overshooting of mean phenotypes regulate densities

Frean and Abraham (2001) claim that the central fixed point is a repeller in well-mixed finite-population RPS because the densities 'overshoot' the stable orbits of the mean field model. In the intraspecific competition model, the finite population also appears to lead to overshooting of mean-field orbits, but it is the traits which overshoot these orbits, not the densities. In the agent-based model, figure 7.5(a) shows that the initial fluctuations in ρ decrease with time while the fluctuations in \bar{x}_A , \bar{x}_B increase with time. In the intraspecific competition model, overshooting in trait values *away from* the fixed point leads to an overcorrecting effect on the densities *towards* the fixed point. This is the case in spite of the fact that densities are able to change faster than the trait values (as long as $\mu < 1$).

7.3.2 Mean traits versus trait distribution

A second reason for the disparity in the behaviour of the infinite and finite population models is that an entire distribution of trait values is represented explicitly in the finite population model, while the infinite population model only represents the mean trait value.

In the finite population case, the distribution quickly spreads out from the initial mean to cover a wide spectrum of possible phenotypes, shown in figures 7.5(e) and (f), until eventually all the possible phenotypes are present at least in small numbers, as shown in the final trait distributions in figure 7.1. Once this wide genetic variation is available, the system is able to dampen large changes in densities very quickly because the necessary phenotypes no longer need to evolve; they are already present in the population.

The infinite population system used the AD assumption that mutants are rare in equations (7.2). This assumption becomes inaccurate as time passes, because the mutants don't stay rare for very long.

126



Figure 7.6: Nonspatial simulations. (a) With mutations ($\sigma = 0.01$), and (b) without mutations. In both cases the population size is 2,250,000, $\beta = 1$, and $\rho_0 = 0.5$.

7.3.3 Flat trait distributions in nonspatial simulations

While the numerical model explains why densities are regulated at $\frac{1}{2}$, it cannot explain why the trait distributions eventually flatten out in figure 7.1, because it doesn't represent those distributions at all.

Symmetric competition

When competition is symmetric ($\beta = 1$), the uniformity of the x_i s is due to small mutations slowly spreading out over the trait space. Once the mean traits \bar{x}_A and \bar{x}_B are at $\frac{1}{2}$, single individuals are relatively indifferent about their level of intraspecific competitiveness, because the other agents they meet are effectively the average ones, so there is no force pushing trait values in either direction.

The role of drift in flattening the distributions can be confirmed by turning off mutations. Figure 7.6 shows two simulations, both of which start with an initial non-uniform distribution with $p(x_i)$ s centred on $\frac{1}{2}$. The x_i distribution becomes uniform in the presence of mutations in figure 7.6(a), but remains relatively static in their absence in figure 7.6(b), indicating that selection is



Figure 7.7: Average over 20 nonspatial simulation runs of distributions of x_i s after 15,000 generations with 1,000,000 individuals for β in (0.25, 0.35, 0.45 0.55, 0.65, 0.75), $\rho_i, \sigma, p(x_A), p(x_B)$ as in fig. 7.1. The dotted lines indicate the mean value of x for the species in the plot.

very weak.

Asymmetric competition

When competition is asymmetric, for example in figure 7.1(b) in which $\beta = 0.6$, mean levels of intraspecific competition tend to move towards the values which balance out the innate mismatch in competitiveness between the two species, as this is what keeps the densities stable at $\rho = \frac{1}{2}$.

With the infinite population model, the fixed point for the intraspecific competitiveness trait is at $1 - f(\bar{z}_A) = \beta(1 - f(\bar{z}_B))$. The corresponding relationship,

$$\frac{1-\bar{x}_A}{1-\bar{x}_B} = \beta,\tag{7.5}$$

holds at the equilibrium in the finite population model as long as neither species goes extinct. The species compete interspecifically exactly as much as they need to in order to balance the difference in their relative strengths.

For example, the equality (7.5) is approximately true for all the simulations in figure 7.7, where $\beta \ge 0.25$. (Extinctions are quite common in the nonspatial system when $\beta < 0.25$). As in the $\beta = 1$ case, mutations flatten out the distribution — figure 7.7 shows results that are nearly the flattest distributions which keep the mean *x*s at their stable values.



Figure 7.8: Simulations with global interactions on a 1000×1000 grid leading to stable coexistence from the starting conditions of scenarios 1 and 2 in section 6.4.1. The solid lines show the densities of the two species, and the dotted lines their mean levels of intraspecific competitiveness. The scenario 1 simulations start with one species *A* individual with x = 0.01, and 999,999 species *B* individuals with x = 0.01. The scenario 2 simulations start with one species *B* individual with x = 0.01 and 999,999 species *B* individual with x = 0.01 and 999,999 species *A* individuals with x = 0.01 and 999,999 species *A* individuals with x = 0.99.

7.4 Importance of spatial structure for coexistence

Because the spatially unstructured system produces stable densities at $\rho = \frac{1}{2}$, it is interesting to examine the effect of long-range dispersal on the intraspecific competition system.

Figure 7.8 shows the densities and mean trait values for some examples which do not end in extinctions, and should be compared with the spatially structured versions in figure 6.9 on page 109.



Figure 7.9: Regions of extinctions and coexistence in the β - σ parameter space for simulations with global interactions starting with the conditions of scenarios 1 and 2. Each point represents the state of a population of size 90,000 after 10,000 generations. If the simulation ended with the extinction of one species, the winning species is marked by a either dark blue point (for species *B*) or a brown point (for species *A*). All other colours from light blue through to white and then orange indicate the species densities after 10,000 generations.

When interactions are global, there are large oscillations in species densities, which occur on a much faster timescale than the spatially structured system, and which continue for much longer. Mean trait levels only respond to changes in densities when those densities cross the line $\rho = \frac{1}{2}$, but they respond immediately rather than with a time lag as in the spatial case.

If the system avoids early extinctions, the trait values will eventually reach the levels that bring ρ to $\frac{1}{2}$ and the system becomes stable. However, as the graphs show, the species densities can get very close to zero on the way, so the likelihood of early extinctions is much higher when interactions are global. Figure 7.9 shows the species densities for finite population simulations of non-spatially structured populations. The only difference between these simulations and those with local interactions in figure 6.10 on page 111 is that the secondary wave of extinctions (see section 6.4.2) is more likely when interactions are global, implying that a well-mixed population will need to have a slightly faster rate σ of phenotypic change in order to stabilise.

7.5 Summary

In the RPS game, stability is due to a spatially-structured population, but the intraspecific competition model is inherently stable in the absence of spatial effects.

An analysis of an infinite-population, mean-field approximation to the intraspecific competition model shows that mean values of the intraspecific competitiveness trait act immediately to repress density fluctuations whenever the densities move away from $\rho = \frac{1}{2}$.

This leads to stable orbits in the infinite-population model, but in a finite population, perhaps because more phenotypic variation is available, the traits overshoot the orbits, pushing the densities towards $\frac{1}{2}$.

In the scenarios described in section 6.4, a well-mixed population will spend more time with densities close to zero, and consequently such a system is more likely to collapse than its spatially structured equivalent when the immigrant is invading at a fast speed and the mutation rate is low.
Chapter 8

Disruptive selection in the intraspecific competition model

In the intraspecific competition model, the individual-level trait is often subjected to disruptive selection in many examples. As a result, levels of intraspecific competitiveness tend towards zero or one while intermediate phenotypes are selected against.

Figure 7.1 on page 118 proves that this divergence is dependent on spatial structure and local interactions, because when spatial structure is removed from the population, individual traits no longer diverge, but smoothly distributed across the trait space.

In this section I try to show why this is the case for the situation in which $\beta = 1$, arguing the following points:

- (1) A small amount of local clustering of individuals of the same species (and with similar x_i s) must necessarily occur because of the nature of the system's invasion dynamics, so occupants of neighbouring cells are more similar on average than occupants of random cells;
- (2) Therefore, a phenotype that is particularly well adapted to colonising one neighbouring cell is also likely to be better adapted than average at colonising that cell's neighbours; and
- (3) Individuals with x_i close to zero or one are the fastest invaders in a clustered environment.

8.1 Spatial interactions create local clustering

It is a necessary consequence of the model's dynamics (sections 6.2.2 and 6.2.3, page 93) that neighbouring cells are more likely to contain similar individuals than cells selected at random.

Because invasions can only happen between neighbouring cells, and because children have the same species and similar value of x to their parent, once there has been a single successful invasion, there will be a single cluster of two similar cells.

Over time, these invasions increase the clustering. In a population of N individuals there are N interactions in one generation, but some of these interactions will involve the same potential invasion victim. The expected number of unique items m' in m draws from a list of M distinct items is $E(m') = M(1 - (1 - 1/M)^m)$ (Cohn, 1992), so the expected number of unique victims in a generation N_u is $N(1 - (1 - 1/N)^N)$. This number is approximately $0.63 \times N$ for reasonable values of N greater than about 100.

The consequence of this is that even if all individuals have the same invasion probability p against every other individual, after one generation there will have been around pN_u successful invasions. So as long as the mean invasion probability $p > 1/N_u$ there will be related neighbours in the population.

8.2 The three-cell model

Once it is accepted that neighbours in the population are more likely to be similar to one another than non-neighbours, a necessary consequence that individuals with extreme levels of intraspecific competition will do better in the short term than those with intermediate levels, because the extreme-valued individuals are faster invaders.

8.2.1 A measure of short-term invasion fitness

To why this happens on a very small scale, I consider three neighbouring individuals, a, i, and j, of species s_a, s_i, s_j and with intraspecific competitiveness levels x_a, x_i, x_j . I assume that a and i are neighbours and that i and j are neighbours, but a and j are not neighbours:

```
a i j
```

The intuition which motivates the argument is that when there is local clustering, if *a* successfully invades *i*, then *a*'s probability of subsequently invading *j* is greater than would be expected without any local clustering, because *i* and *j* are likely to be similar. Being well-adapted to invading any one neighbour will on average result in fractionally more than one successful invasion.

As an approximation of *a*'s short-term invasion fitness, I consider the probability that the third (rightmost) cell is occupied by a type *a* individual after two interactions: an interaction between the left pair of cells followed by an interaction between the right pair of cells. I write *aij* as shorthand for the initial state shown above, and I will call the probability that the rightmost cell is occupied by *a* after two interactions $p_{1,2}(aij \rightarrow aaa)$. This quantity is *a*'s two-interaction invasion fitness. In the following paragraphs I evaluate $p_{1,2}(aij \rightarrow aaa)$ for different amounts of local clustering and show how it varies with x_a .

8.2.2 Quantification of clustering

Local interactions will produce clustering of individuals of the same species, and also clustering of individuals with similar levels of intraspecific competition, so clustering needs to be specified by two numbers. The first, p_s , is the chance that two neighbouring individuals are the same species. If all individuals are equally likely to be of either species and if there are no local interactions, then $p_s = \frac{1}{2}$, because there are two species. If there is local clustering of conspecifics then $p_s > \frac{1}{2}$.

The second measure of clustering, r, can be thought of as the average relatedness of two neighbouring individuals who happen to be of the same species. In genetic terms it is the chance that a single allele in an individual k (which contributes to k's polygenic trait x_k) is also shared by k's conspecific neighbour^{*}.

For the purposes of this discussion I use *r* to describe the chance that any two neighbours *i* and *j* are related given $s_i = s_j$, where 'related' means that they share a recent common ancestor and so $x_i \approx x_j$. Without any local clustering, *r* will be close to zero.

Because mutations are so small, I will assume that in the initial few rounds

^{*}Knowing that two individuals are related entitles us to infer that their phenotypes are similar, but knowing that two individuals have similar phenotypes does not imply that they are genetically related, because it is assumed that there are many different alleles in the population that have similar phenotypic effects.

of a simulation, neighbouring *i* and *j* are *identical* with probability $p_s r$, even though in fact they will just be very similar.

8.2.3 Neighbour relationships

When describing interactions there are three possible kinds of relationships between any two individuals that must be considered in order to determine the chance that an individual invades its neighbour: The two neighbours can be either

- (1) the same species and closely related (with probability $p_s r$), or
- (2) the same species but unrelated (probability $p_s(1-r)$), or
- (3) different species (probability $1 p_s$), which also implies that they are unrelated, since children always inherit their parent's species.

8.2.4 Possible phenotypes in the initial state

Because I have assumed that related individuals are identical, all the possible individuals occupying the initial three cells will be simplified into three categories, where $s_a = A$, and the other species is *B*:

- (1) Phenotypes identical to *a*,
- (2) Random (unrelated to *a*) phenotypes of species *A*, which I will call *a*', *a*'', etc., and
- (3) Random species *B* phenotypes, which I will call b', b'', etc.

8.2.5 Nine initial states

In the initial state the leftmost cell is always occupied by a, and there are two relationships to consider, the relationship between a and i, and the relationship between i and j. Because each of these two relationships can be in one of the three states listed in section 8.2.3, there are nine different initial states, shown in table 8.1, in which the contents of cells i and j are either replaced by a copy of a, or by a random phenotype from the set $\{a', a'', b', b''\}$.

I assume here that s_a and s_j are independent given s_i , and that x_a and x_j are

Relationships			Stat	e	State
Species	Traits	a	i	j	frequency
$s_a = s_i = s_j$	$x_a = x_i = x_j$	a	a	a	p _s rp _s r
$s_a = s_i = s_j$	$x_a = x_i$	a	a	a'	$p_s r p_s (1-r)$
$s_a = s_i \neq s_j$	$x_a = x_i$	a	a	b'	$p_s r(1-p_s)$
$s_a = s_i = s_j$	$x_i = x_j$	a	a'	a'	$p_s(1-r)p_sr$
$s_a = s_i = s_j$		a	a'	a''	$p_s(1-r)p_s(1-r)$
$s_a = s_i \neq s_j$		a	a'	b'	$p_s(1-r)(1-p_s)$
$s_a \neq s_i = s_j$	$x_i = x_j$	a	b'	b'	$(1-p_s)p_sr$
$s_a \neq s_i = s_j$		a	b'	b''	$(1-p_s)p_s(1-r)$
$s_i \neq s_j = s_a$		a	b'	a'	$(1-p_s)(1-p_s)$

Table 8.1: The nine initial states and their frequencies.

independent given x_i .*

8.2.6 Approximation of the system dynamics

The short term invasion fitness of a, $p_{1,2}(aij \rightarrow aaa)$, (over two interactions) is the probability that the rightmost cell j is occupied by a after an interaction between the left pair of cells followed by another interaction between the right pair. This two-interaction fitness quantity is described in terms of invasion probabilities of single interactions.

To model a single interaction between any two neighbours k and l, I assume there is a probability of $\frac{1}{2}$ that k has the opportunity to invade l first, and that if k fails to invade, then l has an opportunity to invade k. I am concentrating on the $\beta = 1$ case where both species are equally competitive, so the values of x_k and x_l alone determine the success of invasions. Because x is an individual's competitiveness with a member of the same species, and 1 - x an individual's ual's competitiveness with a member of the other species, then $p(k \rightarrow l)$, the

^{*}In fact, in a real grid, if there is any local clustering produced by the system, then there is quite likely to be some amount of relatedness between *a* and *j* independent of *i*. However, because the relatedness of individuals that are two or more cells distant is underestimated here, this assumption will not weaken the argument.

probability that *k* wins the contest between *k* and *l*, and invades *l*'s cell, is

$$p(k \to l) = \frac{1}{2}x_k + \frac{1}{2}(1 - x_l)x_k$$

= $x_k(1 - \frac{1}{2}x_l)$ when $s_k = s_l$;
 $p(k \to l) = \frac{1}{2}(1 - x_k) + \frac{1}{2}(1 - (1 - x_l))(1 - x_k)$
= $(1 - x_k)\frac{1}{2}(1 + x_l)$ when $s_k \neq s_l$.

These dynamics are different to the way the simulations in section 6.2 are run. In the simulation, the potential killer has a chance to invade the potential victim, and that is the end of the interaction. In this model, the two individuals have equal chances at getting the first attack. If the first attacker fails, the attacked individual has an opportunity to invade the first attacker. This change has been made here to simplify the system so that only two interactions need to be analysed, while retaining the possibility of invasions occurring in either direction.

8.2.7 Interaction fitness and invasion probabilities

The invasion probabilities above are being used here as part of a measure of *a*'s fitness, but they are not exactly the same thing as *a*'s fitness. The benefit to the copy of *a* in the leftmost cell of copies of *a* in the rightmost cell is the same whether or not the leftmost copy successfully invades the rightmost one.

 $p_1(aij \rightarrow aaj)$ describes the probability that the second cell is occupied by a copy of *a* after the first interaction, and $p_2(aaj \rightarrow aaa)$ the probability that the third cell is occupied by a copy of *a* after the second interaction.

When $i \neq a$, $p_1(aij \rightarrow aaj)$ is the same as *a*'s invasion probability $p(a \rightarrow i)$ described above. When i = a, *a*'s expected fitness is always one because *a* is indifferent as to whether or not she successfully invades a copy of herself. The interaction fitness of *a* is defined by

$$\begin{array}{ll} p_1(aaj \rightarrow aaj) = 1, \\ p_1(aij \rightarrow aaj) = p(a \rightarrow i), & i \neq a, \\ p_2(aaa \rightarrow aaa) = 1, \\ p_2(aaj \rightarrow aaa) = p(a \rightarrow j), & j \neq a. \end{array}$$

For the nine initial states, *a*'s single interaction fitness is listed in the second and fourth columns of table 8.2.

i	$p_1(aij \rightarrow aaj)$	j	$p_2(aaj \rightarrow aaa)$	mean $p_{1,2}(aij \rightarrow aaa)$
а	1	a	1	1
a	1	a'	$x_a(1-\frac{1}{2}x_{a'})$	$\frac{3}{4}x_a$
а	1	b'	$(1-x_a)\frac{1}{2}(1+x_{b'})$	$\frac{3}{4}(1-x_a)$
a'	$x_a(1-\frac{1}{2}x_{a'})$	a'	$x_a(1-\frac{1}{2}x_{a'})$	$\frac{7}{12}x_{a}^{2}$
a'	$x_a(1-rac{1}{2}x_{a'})$	a''	$x_a(1-\tfrac{1}{2}x_{a''})$	$\frac{9}{16}x_{a}^{2}$
a'	$x_a(1-rac{1}{2}x_{a'})$	b'	$(1-x_a)\frac{1}{2}(1+x_{b'})$	$\frac{9}{16}x_a(1-x_a)$
b'	$(1-x_a)\frac{1}{2}(1+x_{b'})$	b'	$(1-x_a)\frac{1}{2}(1+x_{b'})$	$\frac{7}{12}(1-x_a)^2$
b'	$(1-x_a)\frac{1}{2}(1+x_{b'})$	b''	$(1-x_a)\frac{1}{2}(1+x_{b''})$	$\frac{9}{16}(1-x_a)^2$
<i>b</i> ′	$(1-x_a)\frac{1}{2}(1+x_{b'})$	a'	$x_a(1-\frac{1}{2}x_{a'})$	$\frac{9}{16}x_a(1-x_a)$

Table 8.2: Frequencies and invasion probabilities for each of nine possible relationships between three neighbouring cells.

8.2.8 Expected two-interaction fitness

a's fitness after two interactions, $p_{1,2}(aij \rightarrow aaa)$, is just the product *a*'s fitness from the two single interactions:

$$p_{1,2}(aij \rightarrow aaa) = p_1(aij \rightarrow aaj) \times p_2(aaj \rightarrow aaa),$$

and this is calculated for each of the nine initial states in table 8.2. For example in the case on row 5 of the table, where i and j are both of species A but all three individuals are unrelated,

$$p_{1,2}(aa'a'' \to aaa) = p_1(aa'a'' \to aaa'') p_2(aaa'' \to aaa) = x_a^2(1 - \frac{1}{2}x_{a'})(1 - \frac{1}{2}x_{a''}).$$

Nothing is known about $x_{a'}$ and $x_{a''}$ here, because a' and a'' are completely unrelated to a, and to one another. However, if it is assumed that $x_{a'}$ and $x_{a''}$ are uniformly distributed in [0,1], then a's expected fitness can be found by integrating:

$$E(p_{1,2}(aa'a'' \rightarrow aaa) = \iint_0^1 p_1(aa'a'' \rightarrow aaa'') p_2(aaa'' \rightarrow aaa) dx_{a'} dx_{a''}$$
$$= \frac{9}{16}x_a^2.$$

Expected fitnesses can be expressed for all nine states by integrating out all occurrences of $x_{a'}$, $x_{a''}$, $x_{b'}$, and $x_{b''}$. The resulting mean fitnesses are listed in the last column of table 8.2.

8.2.9 Total expected short-term invasion fitness

The total expected fitness of *a* over two interactions is obtained by summing over the mean values of $p_{1,2}(aij \rightarrow aaa)$ for each of the nine states in the last column of table 8.2, weighted by the frequencies of the nine states given in the second column of table 8.1:

$$p_{1,2}(aij \to aaa) = p_s r p_s r + p_s r p_s (1-r) \frac{3}{4} x_a + p_s r (1-p_s) \frac{3}{4} (1-x_a) + p_s (1-r) p_s r \frac{7}{12} x_a^2 + p_s (1-r) p_s (1-r) \frac{9}{16} x_a^2 + p_s (1-r) (1-p_s) \frac{9}{16} x_a (1-x_a) + (1-p_s) p_s r \frac{7}{12} (1-x_a)^2 + (1-p_s) (1-r) \frac{9}{16} (1-x_a).$$
(8.1)

8.2.10 Invasion fitness in clustered environments

The mean invasion fitness of equation (8.1) is invariant with respect to x_a as long as there is no clustering (i.e. when $p_s = \frac{1}{2}$ and r = 0). As soon as there is some clustering, however, $p_{1,2}(aij \rightarrow aaa)$ rises. Figure 8.1(a) shows how the values of p_s and r typically rise at the start of grid simulations without mutations. During the first four generations, p_s rises from 0.5 to 0.67 and rrises from 0 to 0.49. Figure 8.1(b) shows how $p_{1,2}(aij \rightarrow aaa)$ varies with x_a for the values of p_s and r which were measured in the simulation of figure 8.1(a). Invasion fitness as measured by $p_{1,2}(aij \rightarrow aaa)$ is highest when $x_a = 0$ and $x_a = 1$, and at its lowest around $x_a = \frac{1}{2}$.

Although there is no method for estimating the extent of clustering that will occur in any particular SCA simulation, the fact that x_i values do indeed tend towards zero or one indicates that $p_{1,2}(aij \rightarrow aaa)$ is a good measure of invasion fitness.

8.3 The fastest invaders survive

In the long term, it is phenotypes with *x* close to zero and one that will tend to survive in the grid simulations described in chapter 6. It is also these phe-



Figure 8.1: (a) Average values for p_s and r over five simulation runs on a 500 × 500 grid after 0,1,2,3, and 4 generations, without mutations. In the non-spatial case, p_s and r remain at 0.5 and 0 respectively. (b) Individual invasion fitness $p_{1,2}(aij \rightarrow aaa)$ as defined in equation (8.1) for the five values of p_s , r shown in (a).

notypes that have the highest two-interaction invasion probabilities.

In the short term, the success of these fast invaders will tend to reinforce the clustering. On average, they invade new territory slightly faster once there is a small amount of clustering in the system, so slightly larger areas will be occupied by these fast invaders. These areas will then be more likely to be colonised first by other individuals with the best-adapted x to those clusters, and the ones with the best-adapted x will themselves be fast-invading extreme-x individuals.

Intermediate-*x* individuals will on average succumb to invasion more slowly, but they cannot take advantage of the clustering as well as the faster invaders. Any slightly faster mutants of these slower individuals will face a short-term advantage and intermediate phenotypes will eventually lose out.

8.4 Asymmetric competition

The above argument only applies to the case in which both species are equally matched ($\beta = 1$), but the general principle should apply when $\beta < 1$. The preconditions for extreme divergence in x_i s also apply when competition is asymmetric: clustering will still be prevalent in the grid, and individuals with x close to zero or one are still expected to be the fastest invaders. However, these potential fast invaders will only be the actual fast invaders if there is someone for them to invade. In the competition between these fast, extreme phenotypes, the identities of the most likely survivors were established in section 6.3.4 on page 103.

8.5 Summary

In the basic case where $\beta = 1$, the divergence in *x* values occurs only when the interactions in the system are local, as shown in figure 7.1 (page 118), so it is spatial structure that is responsible for the observed specialisation in interor intraspecific competitiveness.

The profitability of specialisation can be explained with the definition of shortterm invasion fitness, which shows that the extreme-x phenotypes have a slight advantage when looking ahead only two grid interactions. Their advantage comes about because local interactions cause spatial clustering of similar phenotypes. An individual that successfully invades one victim is statistically more likely than average to also be successful at invading that victim's neighbour, because the neighbour is more likely than average to be similar to the first victim. The result is large clusters of specialists, who are most likely to be invaded first by other specialists. All this specialisation, however, also makes individuals vulnerable, and a constantly shifting landscape is the result.

144 CHAPTER 8. DISRUPTIVE SELECTION IN THE INTRA MODEL

Chapter 9

Cyclic competition among toxin-producers

The intraspecific competition model described in chapter 6 requires narrowly targeted species-specific forms of competition. This restriction limits its potential usefulness, and motivates the question of whether there are ecosystems that work in this way.

It turns out that in ecosystems where toxin-based competition is prevalent, there is a close correlation with the model's assumptions, for two reasons. Firstly, toxins can be produced in varying quantities, so the offensive ability of a toxin-producer is likely to be able to vary from individual to individual rather than being completely species-dependent. Secondly, toxins are often effective against a large class of organisms, sometimes against entire species or even families of species, while individuals outside the class are often completely immune to the toxin's effects.

In this chapter I show how intraspecific competition can lead to RPS-like competition in an ecosystem of interacting toxin-producers. In the next section I describe some of the features of ecosystems made up of toxin-producing microbes, and review several recent theoretical models of these kinds of ecosystems. I then show in section 9.1.3 that the model presented in the last chapter is not exactly applicable to known systems of toxin-producers in nature. Section 9.2 defines a variant of the model which is a better fit to real ecosystems, and sections 9.3–9.5 describe the behaviour of this variant, and show that it produces an amazing diversity of monomorphic and polymorphic equilibria, including many involving cyclic competition. Section 9.6 outlines the behaviour of a more general class of models of which the section 9.2 model is a part. Finally, section 9.7 discusses some implications of the behaviour of the models.

9.1 Modelling toxin-based competition

Microbial ecosystems are diverse. Cells of many different species or strains interact in a complex web of competitive relationships, and both inter- and intraspecific competition are intense. Intransitivity between species has been observed, including spatial RPS games among bacterial strains (Kerr et al., 2002).

It has been claimed that the diversity of microbial ecosystems is primarily a result of the production of specific toxins (Pagie and Hogeweg, 1999; Czárán et al., 2002). In bacteria for example, toxins called colicins are produced by strains of *E. coli* which have been shown to inhibit or kill some other strains while leaving still others unaffected (Reeves, 1972). Chemical warfare of a similar nature is widespread among microorganisms in general (Riley, 1998), and also common among the coral reef-dwelling organisms such as corals, algae, and sponges (Bakus et al., 1986; Ates, 1989), organisms which are known to compete for space in nontransitive relationships (Buss and Jackson, 1979).

9.1.1 Theoretical models of microbial toxin systems

Because of the ease with which bacteria can be studied in the laboratory, the nature of some microbial toxins is quite well understood, and this has led to a number of theoretical models of toxin-based competition. The primary aim of most of these models is to explain why the observed diversity of the ecosystems is so high. The following paragraphs summarise some of these results before comparing the models with the intraspecific competition model from chapter 6.

Frank (1994)

Frank presents a reaction-diffusion model of colicinogenic bacteria which describes competition for scarce resources between a colicin-producing strain and a sensitive strain. One of the two strains wins out when there is no spatial structure to the environment, and a stable polymorphism between the two types results when there is variation in the habitat quality of different spatial locations.

9.1. MODELLING TOXIN-BASED COMPETITION

The model does not produce stable polymorphisms due to cyclic competition, but establishes the importance of a spatial population structure.

Durrett and Levin (1997)

Durrett and Levin compare colicinogenic and sensitive bacterial strains in spatially structured and unstructured environments using stochastic cellular automata. The model is complex, assigning parameters for birth rate, death rate, and colicin-producing rate to each species.

In the well-mixed environment, one of the two strains goes extinct depending only on their initial densities. In the spatially structured environment, one of the strains also goes extinct, but in this case which of the two strains goes extinct depends on the metabolic cost of the toxin production, and is independent of the initial densities.

When the competition occurs between a sensitive type, a high-cost colicinproducer, and a third 'cheater' type which produces colicin at a slower rate, the three types will sometimes end up coexisting in a rock-paper-scissors configuration, in which the colicin-producer beats the sensitive type, which beats the cheater, which in turn beats the colicin-producer.

Iwasa et al. (1998)

Iwasa et al. analyse a two-species colicin system very much like Durrett and Levin's using pair approximations, and find that for some regions of the parameter space, the winner depends on its strengths as determined by the parameters, but that for other areas of parameter space the winner depends on the initial densities. The pair approximation appears to give a result midway between the completely nonspatial situation, in which initial densities almost always determine the winner, and the explicit spatial models, in which they almost never do.

Gordon and Riley (1999)

Gordon and Riley provide empirical results from experiments in which several different strains of colicinogenic bacteria each compete with a sensitive strain, and compare these results with those of a mathematical model which makes realistic assumptions about the invasion dynamics of the cells in the experimental setup. Colicin-producing cells are added to a sensitive population at regular intervals, in a well-mixed environment, and the additions were continued until the sensitive cells were no longer present.

148 CHAPTER 9. CYCLIC COMPETITION AMONG TOXIN-PRODUCERS

The speed with which colicin-producing strains invade sensitive strains increased with the initial density of colicin-producing strain, and also increased with the strength of the colicin.

Pagie and Hogeweg (1999)

Pagie and Hogeweg model a colicinogenic bacterial ecosystem with a large number of different toxins present, in which each individual can produce, and acquire immunity to none, some, or all of these toxins. The resulting systems are very diverse, and after time the system falls into one of two states: one called *hyperimmunity mode*, where bacteria produce few toxins but are immune to most toxins; and another called *multitoxicity mode*, where many toxins are produced but bacteria are not generally immune to toxins which they don't produce themselves. Which of these two modes the system falls into depends on the growth penalty which the bacteria pay for being immune. If the immunity penalty is low, the system tends towards hyperimmunity; if the penalty is high, the system tends towards multitoxicity.

Czárán et al. (2002)

Czárán et al. present a model similar to that of Pagie and Hogeweg (1999), but which they claim is intended to generalise over all microbial systems rather than specifically applying to colicin-producing bacteria. There are 14 possible toxins and 14 corresponding immunities, where the penalties are restricted so that for each toxin there is a rock-paper-scissors like payoff structure as suggested in Durrett and Levin (1997), in which killer, sensitive, and resistant strains compete cyclically.

The hyperimmunity mode found by Pagie and Hogeweg (1999) is reproduced in the model, but the multitoxicity mode has 'frozen' spatial dynamics with little change over time, unlike Pagie and Hogeweg's multitoxicity, which has a very high turnover of toxin frequencies.

Czárán et al. also claim that different interpretations and parameter values in the model apply to different kinds of ecosystems. For bacteria, it makes sense to think of each strain as a variant within a single species, and in this case the rate of acquisition of novel toxins is high, because of the high frequency of horizontal transfer of colicin-encoding plasmids between strains of bacteria. For yeasts, however, each 'strain' in the model represents a species, and the rate of acquisition of novel toxins by a species is low.

9.1.2 Advantages of the intraspecific model

There are three aspects of toxin-based competition which previous work has not brought together which are able to be combined in a version of the intraspecific competition model.

Degrees of toxin production

With the exception of Gordon and Riley (1999), none of these theoretical studies examine what happens when small changes are allowed to the amount or the strength of toxin produced. This is despite the fact that this sort of variation is known to exist (see Riley and Gordon (1996)). The majority assume all-or-nothing toxin production because it means the models are more easily analysed. For example, Iwasa et al. (1998) use pair approximations, a technique which is only possible when there are a small number of phenotypes to consider. Durrett and Levin (1997) consider RPS competition between strains with two discrete levels of colicin-production (the colicin-producer and the cheater) but do not consider evolution in these colicin levels.

Local spatial interactions

Only some of the previous models consider the effects of spatial interactions, even though they are known to be extremely important in bacterial ecosystems. For example, Kerr et al. (2002) show that they can be the difference between a stable RPS community and a monoculture. In the only study I am aware of which considers the evolution of degrees of toxin production Gordon and Riley (1999), a well-mixed environment is assumed.

Generality

Some of the models include aspects of the toxic competition process that are specific to colicin producing bacteria, rather than general forms of toxin-based competition. For example, Gordon and Riley (1999) explicitly model the concentration of colicin in the environment, and the fact that colicin-producing cells must undergo lysis before the colicin is released. On the other hand, Pagie and Hogeweg (1999) and Czárán et al. (2002) provide models which do not rely on specific features of bacterial colicin production or features of a particular experimental setup.

The intraspecific model can combine smooth variation in toxin levels with local spatial effects, and is also generalisable to different types of toxin pro-

duction. It is potentially valuable because it may explain why all-or-nothing toxin production should be expected in spatial toxin systems, even when it is not initially present.

9.1.3 Application of the intraspecific model

The intraspecific competition model can only represent species that produce *inter*specific toxins. The two species, *A* and *B*, are both capable of producing a toxin that kills the other at a metabolic cost. All *A*s are presumed to be sensitive to the toxin produced by *B*s, and all *B*s are presumed to be sensitive to the toxin produced by *A*s. The amount of toxin produced by *i* is represented by $1 - x_i$ in the model, because the amount of toxin produced is *i*'s interspecific competitiveness. The production of low levels of the toxin by *i* (high x_i) makes *i* more competitive against its own type because only a small cost is paid, so it is able to reproduce more quickly. But when an individual fails to produce the toxin, it becomes vulnerable to the toxin produced by the other type.

Prediction of the intraspecific competition model

In the instance that two strains of bacteria with the qualities assumed in this section could be isolated and allowed to interact in the lab, the model of the last section would predict the following:

- (1) Eventually all the bacteria would be producing either some maximum amount of the toxin or no toxin at all.
- (2) If the two original strains grow at moderately different rates, then the model predicts that eventually all of the slower-growing strain will be producing the maximum amount of toxin, while individual cells of the faster-growing strain will produce either the maximum amount or none at all.

Interspecific toxins

It has often been assumed that colicin production is a form of intraspecific competition, and that colicins only have an effect on individuals that are closely related to the producer. However recent investigations show that colicins also play a role in interspecific competition, and that some of the colicins produced by *E. coli* may affect other less closely related species of bacteria (Riley, 1998).

Horizontal transfer between strains

The likelihood of horizontal transfer of genes that encode for toxin production or toxin resistance is an important factor. The model will only apply when horizontal transfer is unlikely in the timescale under consideration, otherwise, morphs which produce, or are resistant to both toxins may arise. In these cases the resulting system dynamics will be more like those of Pagie and Hogeweg (1999) and Czárán et al. (2002) in which selection favours morphs with multiple resistance or multiple toxicity.

There is evidence that although horizontal transfer of toxin-encoding genes is common in bacteria, it is rare among yeasts (Wickner, 1992; Abranches et al., 1997), in which case yeast toxins will be interspecific. In any case there are many other types of bacteriocins apart from colicins, antimicrobial substances are also produced by fungi (Starmer et al., 1987; Berdy, 1974), and as previously mentioned, interspecific toxins are known among reef-dwellers.

Tradeoff between growth rate and toxin production

Adams et al. (1979) have shown that a trade-off exists between the production of colicin and an organism's growth rate: colicin-producing strains have a slower growth rate than sensitive strains. It is highly likely that most toxins are expensive to produce, so it is important that the model reproduce this feature.

While the model has a tradeoff between intraspecific competition (no toxin production, x = 1) and interspecific competition (toxin production, x = 0), this is not the same as a tradeoff between toxin production and growth rate.

The model as it stands assumes that *i*'s level toxin production is represented by $1 - x_i$, in which case x_i should represent *i*'s growth rate. But growth rate should apply equally against all others regardless of their species, and the model's assumptions entail that x_i only applies against members of the same species. Therefore, a similar model in which growth and toxicity explicitly trade off against one another is explored in the rest of this chapter.

9.2 A two-species toxin-growth tradeoff model

In this model, individuals of two species interact on a grid, in a similar way to that of the intraspecific competition model from section 6.2 (page 92). How-

152 CHAPTER 9. CYCLIC COMPETITION AMONG TOXIN-PRODUCERS

ever, instead of facing a tradeoff between interspecific and intraspecific competition, here individuals trade interspecific toxin production against their growth rates, where growth rate is form of competition that affects both species equally and which is less effective than toxin-production.

Individually-determined toxin-production trait

Every individual *i* varies in the amount of toxin it produces, and *i*'s toxin production level will be described by v_i , where $0 \le v_i \le 1$. Because growth rate trades off against toxin production, *i*'s growth rate is described by $1 - v_i$.

Species-determined competitiveness

The basic competitiveness of two species numbered 0 and 1 is divided into species-determined growth rates g_0 and g_1 , and species-determined toxicities h_0 and h_1 .

A species *s* individual *i*'s total competitiveness is found by multiplying its individual growth rate $1 - v_i$ by its species-determined growth rate g_s , and this total growth rate $(1 - v_i)g_s$ is used in competition against all other individuals regardless of their species.

A species *s* individual *j*'s total poisoning rate is its individual rate of toxin production v_j multiplied by its species-determined toxin strength h_s . In contrast to individual growth rate, the total individual poisoning rate v_ih_s is effective only against individuals of the other species, because the model assumes that all individuals are immune to the toxin which is produced by their own species.

Invasion probabilities

The probability $p_{i \rightarrow j}$, that individual *i* invades individual *j* depends on the individuals' species, and is shown in table 9.1(a).

Ratios of species strength and toxin strength

In the intraspecific competition model, the quantity $\beta = c_b/c_a$ described the competitiveness of the weaker species *B* relative to the stronger species *A*, but here species *A* and *B* compete in two ways; they have inherent growth rates g_A and g_B , and inherent toxin strengths h_A and h_B .

Initially I will assume that the 'weakness ratio' of the weaker species, β , stems

s _i	s _j	$p_{i \rightarrow j}$		s _i	s _j	$p_{i ightarrow j}$
0	0	$g_0(1-v_i)$		A	A	$\gamma(1-v_i)$
0	1	$g_0(1-v_i)+h_0v_i$		A	В	$\gamma(1-v_i)+v_i$
1	0	$g_1(1-v_i)+h_1v_i$		В	A	$\beta(\gamma(1-v_i)+v_i)$
1	1	$g_1(1-v_i)$		В	В	$eta\gamma(1-v_i)$
(a)			(b)			

Table 9.1: The outcome of a single interaction is determined by the probability p_{ij} that an individual *i* invades a site occupied by another individual *j*.

from some underlying ability of the species to use resources more or less efficiently, and therefore that β applies to both growth and toxin production, so that

$$\beta = \frac{g_B}{g_A} = \frac{h_B}{h_A}.\tag{9.1}$$

I will also assume that releasing toxin is a more effective method of competition than mere growth, and define a second 'weakness ratio' γ of growth relative to poisoning, such that

$$\gamma = \frac{g_A}{h_A} = \frac{g_B}{h_B},\tag{9.2}$$

where $0 < \gamma \leq 1$.

Total individual invasion probabilities are summarised in terms of β and γ in table 9.1(b), where it is assumed that the poisoning rate of the stronger species, h_A , is equal to one.

In fact it is probably unrealistic to assume that both species will produce equally effective toxins, and probably unrealistic to assume that both species' toxins are better than their growth rates by exactly the same amount. Consequently, the effects of relaxing these assumptions about β and γ are examined in section 9.6, although it turns out that many of the model's essential features are captured in the special case where equations (9.1) and (9.2) hold.

9.3 Species densities

Throughout the range of β and γ there are many regions in which both species are viable in the long-term, and many regions with cyclic competition among morphs that diverge in a branching process.

9.3.1 Coexistence and stability

Figure 9.1(a) shows the final species densities and extinctions for simulations on a small grid for β and γ from 0.005 to 1. The stronger species always wins whenever $\beta < \gamma$, and these extinctions are shown as the large grey area in the bottom right of the graph.

Outside this area, extinction of the weaker species is common when γ is close to zero or one, that is, when the toxin is either very weak or very strong.

Extinction of the stronger species is very rare, and occurs only for a small number of simulations shown by the black dots in the picture, in the area where β is only slightly larger than γ .

In the region where coexistence is the usual outcome, there is a critical value of γ , just greater than $\frac{1}{2}$, in which the two species finish with relatively similar densities. This appears as a vertical white strip just to the right of the dotted line marking $\gamma = \frac{1}{2}$. As γ decreases or increases away from this critical value, the equilibrium density of the stronger species gradually increases until the saturation points are reached.

The final species densities in this model stabilise away from $\frac{1}{2}$ in most cases where the two species coexist. This is in marked contrast to the intraspecific competition model, in which coexistence almost always ends with species densities stabilising around $\frac{1}{2}$ (see figure 6.1, page 96). What both models share is the fact that if the species coexist, then the value of β has only a very minor effect on their equilibrium densities.

9.3.2 Stability in spatially unstructured populations

The species densities after 10,000 generations are much less predictable when the population has no local spatial structure, as shown in figure 9.1(b). If both species coexist, lack of spatial structure has little effect on the equilibrium, but there is much more randomness because of large high-frequency fluctuations. Extinctions are more likely, as are simulations in which the weaker species has a greater density after 10,000 generations. This effect is only apparent when σ is fairly low — for example, the simulations in figure 9.1(b) were all performed at $\sigma = 0.003$. When σ is higher, the greater adaptability of individuals tends to stabilise the system.

The fact that the species densities are roughly the same in spatial and nonspatial computer simulations suggests that some aspects of the system could be determined without the need for simulations at all, and this is in fact the



Figure 9.1: Extinction, coexistence, and species density in the toxin vs. growth model described in section 9.2. Each point shows the result of a simulation. Grey areas indicate saturation by the stronger species, black points saturation by the weaker species, and shades from orange through to white and then to blue indicate coexistence after 10,000 generations, where the final densities of the two species are indicated by the colour on the scale at the right. The simulations were carried out on a small 220 × 220 grid with initial species densities at $\frac{1}{2}$, uniform initial distribution of toxin production, and mutation rate $\sigma = 0.003$.



Figure 9.2: Net invasion rates among the extreme individuals in the toxin-growth model where v = 0 or v = 1. A0 and B0 are fast-growing (v = 0) competitors of species A and B respectively, and A1 and B1 are maximum toxin-producing (v = 1) competitors of species A and B respectively. The labels on the arrows are the net invasion rates between these individuals from table 9.1.

case. The rest of this section attempts to explain those aspects of the model's behaviour that do not depend on spatial interactions or finite populations, first of all by considering the net invasion rates among extreme individuals, and secondly by considering a mean-field, infinite-population version of the model.

9.3.3 Net invasion rates for pure strategies

An examination of the net invasion rates among pure strategies (individuals with toxin production at zero or one) reveals that there is a cyclic dominance among the four types as long as $\beta > \gamma$ (see figure 9.2). If $\beta < \gamma$ then the non-toxin producing species *A* type has absolute dominance over all others, so a species *A* monoculture should be expected in this case, and that is in fact the result of the simulations as shown by the grey area in the lower right of both pictures in figure 9.1.

This invasion diagram is more complex than the corresponding one for the the intraspecific competition model (figure 6.6(a) on page 102). There are three cycles in the toxin-growth graph and only two in the intraspecific competition graph. Two of the pure strategies (*B*0 and *A*1) in figure 9.2 can potentially be 'cut out' of the main loop on the outside, whereas only one strategy is vulnerable in figure 6.6(a).



Figure 9.3: Relative net invasion rates among the pure strategies in the toxin-growth model for four combinations of β and γ , showing four instances of the graph in figure 9.2. In each of the four cases, the thinnest arrow represents the slowest net invasion rate, and the widths of the other arrows are scaled relative to the slowest. Low γ and β destabilises the system.

The complexity of the pure strategy graph makes prediction of the vulnerable species very difficult. We can vary the values of β and γ , and see which of the net invasion rates become large (figure 9.3), but because of the multiple intransitivities, it is not clear which strategies will have the lowest densities.

9.3.4 Mean-field model

The fact that densities are equal around $\gamma = \frac{1}{2}$ can be explained in part by looking at a mean-field-type model of the toxin-growth system using a similar method to the one described in section 7.1 (page 118). This is a non-spatial, infinite-population model with an explicit parameter μ that defines the relationship between the evolutionary and ecological timescales.

Applying the method in that section gives the following equations, which describe the change in density of species *A*, and the average toxin-production levels of both species:

$$\frac{d\rho}{dt} = \rho(1-\rho) \left[\gamma(1-\beta) - (\beta \bar{v}_B - \bar{v}_A)(1-\gamma) \right]$$

$$\frac{d\bar{v}_A}{dt} = \mu(1-\gamma-\rho)$$

$$\frac{d\bar{v}_B}{dt} = \mu\beta(\rho-\gamma).$$
(9.3)

Densities at the fixed point

The last equation implies that $\rho = \gamma$ when the system is at a fixed point, because neither β nor μ is zero. The second implies that in the case where $\rho = \gamma$, then they both must be equal to $\frac{1}{2}$ at the fixed point. This gives some indica-



Figure 9.4: The combination of mean species toxin-production \bar{v}_A , \bar{v}_B required when the mean-field system is at its fixed point. \bar{v}_A and \bar{v}_B must be between zero and one, so the fixed point is only reachable when $\beta > \frac{1}{2}$.

tion as to why the densities are close to $\frac{1}{2}$ when γ is around $\frac{1}{2}$ in figure 9.1.

The first equation does not imply unique values for \bar{v}_A and \bar{v}_B at the fixed point but it restricts them to a plane in which

$$\frac{1+\bar{v}_A}{1+\bar{v}_B} = \beta.$$

Possible stable combinations of \bar{v}_A , \bar{v}_B at point where the densities are fixed are shown in figure 9.4 for several values of β . This figure shows that because both \bar{v}_A and \bar{v}_B range from zero to one, β must be bigger than $\frac{1}{2}$ when the densities are stable. This is in agreement with the simulation results, which show densities at $\rho = \frac{1}{2}$ close to the line where $\gamma = \frac{1}{2}$, but only where β is also greater than $\frac{1}{2}$.

Symmetric species

When both species are equally competitive, $\beta = 1$ and equations (9.3) reduce to

$$d\rho/dt = \rho(1-\rho)(\bar{v}_A - \bar{v}_B)(1-\gamma)$$

$$d\bar{v}_A/dt = \mu(1-\gamma-\rho)$$

$$d\bar{v}_B/dt = \mu(\rho-\gamma).$$

When this system is initialised with ρ , \bar{v}_A , and \bar{v}_B all equal to $\frac{1}{2}$, $d\rho/dt$ will be zero and \bar{v}_A , \bar{v}_B will both move in the same direction at the same rate $\mu(\frac{1}{2} - \gamma)$.

When $\gamma \neq \frac{1}{2}$, toxin-production will either rise to its maximum or fall to its minimum. When $\gamma < \frac{1}{2}$ both species tend towards maximum toxin production because the toxin is strong, and when $\gamma > \frac{1}{2}$ both species tend towards no toxin production because the toxin is weak. The densities should remain around $\frac{1}{2}$.

Asymmetric species

When $\beta < 1$, there are no interesting fixed points when $\gamma \neq \frac{1}{2}$ which involve coexistence, but when \bar{v}_A and \bar{v}_B are artificially fixed whenever they reach zero from above or one from below, the system can reach a stable state when $\gamma \neq \frac{1}{2}$.

In all cases when $\beta < 1$, the weaker species will tend to produce more toxin than the stronger species in order to balance out the asymmetry in the species strengths. This implies certain facts about the species densities at the points where \bar{v}_A and \bar{v}_B reach zero and one.

Strong toxins

When the toxins are strong, \bar{v}_B rises to one and stays there, and \bar{v}_A settles at a lower level. With \bar{v}_B fixed at one, ρ and \bar{v}_A follow the system defined by

$$d\rho/dt = \rho(1-\rho)[\gamma - \beta - \bar{v}_A(1-\gamma)]$$

$$d\bar{v}_A/dt = \mu(1-\gamma-\rho).$$

This system reaches a fixed point at

$$ar{v}_A = rac{eta - \gamma}{1 - \gamma}, \qquad
ho = 1 - \gamma.$$

Therefore ρ rises as γ falls from $\frac{1}{2}$ to zero because it takes the value $1 - \gamma$ whenever \bar{v}_B is trying to rise above one. This explains the simulation results shown in figure 9.1 where the stronger species becomes more numerous as γ falls from $\frac{1}{2}$ towards zero.

Weak toxins

When the toxins are weak, \bar{v}_A falls to zero and remains there while \bar{v}_B settles at some higher level. With \bar{v}_A fixed at zero, ρ and \bar{v}_B follow the system defined by

$$\begin{split} &d\rho/dt = \rho(1-\rho)[\gamma(1-\beta) - \beta \bar{v}_B(1-\gamma)] \\ &d\bar{v}_B/dt = \mu\beta(\rho-\gamma). \end{split}$$

This system reaches a fixed point at

$$ar{v}_B = rac{(1-eta)\gamma}{eta(1-\gamma)}, \hspace{0.5cm}
ho = \gamma$$

Therefore ρ rises as γ rises, and ρ settles at the value γ when \bar{v}_A is trying to fall below zero. This is also in broad agreement with the simulation result shown in figure 9.1.

9.3.5 Local clustering of species

Figure 9.1(a) shows equal species densities at equilibrium when γ is slightly higher than $\frac{1}{2}$, the theoretical level at which the the densities should be equal according to the mean-field analysis. The most obvious explanation for the discrepancy is that the higher value in the spatially explicit simulations must have something to do with the local clustering that is always a feature of these systems.

In theory, local clustering of conspecific individuals should mean that individuals will be biased in favour of intraspecific competition when compared to the well-mixed system: less toxin-production (lower average levels of v) should be expected in a clustered system.

But a small bias against toxin-production does not account for the observation in the simulations that there is a stable equilibrium where $\rho = \frac{1}{2}$ at $\gamma > \frac{1}{2}$. In fact, I will argue in the following paragraphs that a small amount of clustering of conspecifics should move the point at which species densities are equal and stable to a location in which γ is slightly *less than* $\frac{1}{2}$.

The mean-field model and local clustering

In chapter 4, the pair approximation was used as a way to introduce some level of local structure information into a numerical model. Although the fully-fledged pair approximation cannot be applied to the toxin-growth model (because individuals have continuously-variable traits), I use some aspects of the pair approximation here in order to describe the clustering of species.

In equations (9.3), the mean levels of toxin-production \bar{v}_A , \bar{v}_B respond to the global species densities as defined by ρ . If there is local clustering, then the *A*s and *B*s will 'experience' slightly different ρ s. On average, *A*s will meet more *A*s, and so \bar{v}_A should respond to the densities as if ρ was slightly *higher* than the global density of *A*s. And on average *B*s will meet more *B*s and so \bar{v}_B should change as if ρ was slightly *lower* than the global average.

Pair densities

Using similar notation from the pair approximations in chapter 4, the densities of pairs will be described using the four variables $\rho_{[AA]}$, $\rho_{[AB]}$, $\rho_{[BA]}$, and $\rho_{[BB]}$. In a perfectly mixed population, these pair densities can be described in terms of singleton densities as

$$\rho_{[AA]} = \rho_A \rho_A, \quad \rho_{[AB]} = \frac{1}{2} \rho_A \rho_B, \\ \rho_{[BA]} = \frac{1}{2} \rho_B \rho_A, \quad \rho_{[BB]} = \rho_B \rho_B.$$

$$(9.4)$$

If the world is slightly more clustered, then $\rho_{[AA]}$ and $\rho_{[BB]}$ will be have slightly greater values than those described in equations 9.4, and $\rho_{[AB]}$ and $\rho_{[BA]}$ will have slightly smaller values.

Clustering of conspecifics

Suppose that in the slightly clustered world some proportion ε of the pairs that would have been heterogeneous in the well-mixed environment become homogeneous pairs, and that these extra homogeneous pairs are distributed equally between [AA]s and [BB]s. ε of the pairs that would have been [AB]s under perfect mixing are [AA] pairs due to the clustering, and ε of the pairs that would have been [BA]s are now [BB]s, so that

$$\rho_{[AA]} = \rho_A \rho_A + \varepsilon \cdot \frac{1}{2} \rho_A \rho_B, \quad \rho_{[AB]} = (1 - \varepsilon) \cdot \frac{1}{2} \rho_A \rho_B, \\ \rho_{[BA]} = (1 - \varepsilon) \cdot \frac{1}{2} \rho_B \rho_A, \qquad \rho_{[BB]} = \rho_B \rho_B + \varepsilon \cdot \frac{1}{2} \rho_B \rho_A.$$

$$(9.5)$$

Under these assumptions the global singleton density ρ_A is the same in both the well-mixed and clustered systems, as is ρ_B .

Change in species densities

To update equations (9.3) to take account of this clustering, first note that only the [*AB*] and [*BA*] pairs are important in determining how the two species interact, so the factor $\rho(1 - \rho)$ in the equation for the change in species densities $d\rho/dt$ must change to reflect the clustering. Equations (9.5) imply that the sum of all these pairs is

$$\rho_{[AB]} + \rho_{[BA]} = (1 - \varepsilon)\rho_A \rho_B = (1 - \varepsilon)\rho(1 - \rho).$$

Change in mean levels of toxin-production

Second, the equations for the response of the mean toxin-production traits, $d\bar{v}_A/dt$ and $d\bar{v}_A/dt$ can be updated to take account of the fact that the two

species experience slightly different densities to the global ones. The 'apparent' densities are the conditional probabilities in the pair approximation, so for example to the *A*s, the density of *A*s looks like $p_{A|A}$, the probability that the neighbour of an *A* is an *A*, and to the *B*s, the density of *A*s looks like $p_{A|B}$, the probability that the neighbour of a *B* is an *A*. As established in section 4.3.4 (page 56), with the pair approximation these conditional probabilities are given by

$$p_{A|A} = rac{
ho_{[AA]}}{
ho_A}, \qquad p_{A|B} = rac{
ho_{[AB]}}{
ho_B}.$$

The rates of change of toxin-production traits can be described under the clustered system by substituting $p_{A|A}$ for ρ in the equation for $d\bar{v}_A/dt$ and by substituting $p_{A|B}$ for ρ in the equation for $d\bar{v}_B/dt$.

Modified mean-field equations

The resulting system is described by the three equations

$$\frac{d\rho}{dt} = (1-\varepsilon)\rho(1-\rho)\left[\gamma(1-\beta) - (\beta\bar{v}_B - \bar{v}_A)(1-\gamma)\right]$$

$$\frac{d\bar{v}_A}{dt} = \mu(1-\gamma-\rho-\varepsilon(1-\rho))$$

$$\frac{d\bar{v}_B}{dt} = \mu\beta(\rho(1-\varepsilon)-\gamma).$$
(9.6)

What do these equations imply for the equilibrium species densities? One difference between equations (9.6) and the well-mixed model described by equations (9.3) is that equations (9.6) have no fixed point when $\varepsilon > 0$ (and when $\varepsilon = 0$ the two systems are identical).

When $\varepsilon > 0$, the second and third equations of (9.6) imply that \bar{v}_A and \bar{v}_B can only both be unchanging when $\rho = \frac{1}{2}$, and at this point γ must be slightly *less* than $\frac{1}{2}$ ($\gamma = \frac{1-\varepsilon}{2}$). This is the opposite result to the observations in the simulation results of figure 9.1(a), where $\rho = \frac{1}{2}$ when γ is slightly *greater* than $\frac{1}{2}$.

Figure 9.5 shows the behaviour of the clustered and unclustered systems in more detail. The graphs in the very top row of the figure show the trajectories of ρ , \bar{v}_A and \bar{v}_B over time without clustering, for seven values of γ . As mentioned in section 9.3.4, when $\gamma < \frac{1}{2}$, \bar{v}_B rises to one, and when $\gamma > \frac{1}{2}$, \bar{v}_A falls to zero. After this the system settles into stable cycles in which the long-run average values of ρ , \bar{v}_A , \bar{v}_B are shown by the small diamonds on the right hand edges of the figures. ρ cycles around $\frac{1}{2}$ when $\gamma = \frac{1}{2}$, and averages higher than $\frac{1}{2}$ as γ moves away from $\frac{1}{2}$ in either direction.



Figure 9.5: Behaviour of the system of equations (9.6) for different amounts of spatial clustering ε , $\beta = 0.75$, and $\mu = 0.1$, started from the point at which ρ , \bar{v}_A , \bar{v}_B are all equal to $\frac{1}{2}$. The three rows of graphs at the top show the species densities ρ (black line) and mean levels of toxin production (\bar{v}_A , orange line, and \bar{v}_B , blue line) over time for seven values of γ . The black, orange and blue diamonds on the right hand sides of these graphs show the long-run average of the fluctuations in the three variables if time were continued past the right hand edge. The cube diagrams show a different view of the two top rows. The trajectories start at the centre of the cube, and in most cases they immediately head towards one of the two back walls of the cube (either \bar{v}_A falls to zero or \bar{v}_B rises to one) before starting regular cyclic fluctuations of the other two variables. However, if the amount of clustering balances the toxin strength, then the fluctuations can be stable without either species' toxin-production reaching an extreme value. This is shown by the green line on the left hand diagram, when $\varepsilon = 0$ and $\gamma = 0.5$, and by the light blue line on the right when $\varepsilon = 0.1$ and $\gamma = 0.45$.

The second and third rows show the behaviour of the system as the amount of clustering (ε) increases. The most stable trajectories where ρ cycles around $\frac{1}{2}$ occur when $\gamma = \frac{1-\varepsilon}{2}$. When γ is smaller than this value, \bar{v}_B rises to one and when γ is greater than this value \bar{v}_A falls to zero. In most cases ρ appears to cycle around a value greater than $\frac{1}{2}$ when $\gamma \neq \frac{1}{2}$, and the long run average of ρ gets higher as γ moves away from $\frac{1-\varepsilon}{2}$ in either direction.

Failure of the species clustering model

In summary, the mean field model appears to account for the general phenomenon in which the density of *A* rises as γ increases or decreases away from γ 's critical value, but it predicts a critical γ at exactly $\frac{1}{2}$ rather than the observed value, which is slightly higher.

Extending the analysis to take account of a small amount of local clustering of species doesn't help; instead it suggests that the critical value of γ should be less than $\frac{1}{2}$ rather than greater than $\frac{1}{2}$.

The extension of the mean field model given here is deficient, however, because while it can model the effects of clustering of conspecifics, it does not model the clustering of individuals with similar trait values.

9.4 Distribution of toxin-producers

In the coexistence region of the $\beta - \gamma$ space the two species organise into many different distributions of toxin-producers and non-toxin-producers. The toxin production trait does not always tend towards extreme values, and only relatively small areas of the space fall into the RPS-like pattern.

Figure 9.6 shows the results of simulations similar to the ones depicted in figure 9.1, but it shows the equilibrium toxin-production distributions rather than the species densities.

9.4.1 Classification of toxin-production distributions

Like the intraspecific competition model, the equilibrium trait distributions in the toxin model fall into several classes, each with different qualitative characteristics. The distributions are classified into several types shown by different colours in figure 9.6, depending on where on the toxin-production axis the



Figure 9.6: Overview of the regions of β - γ -space which have qualitatively different outcomes for the toxin-growth model. Each point represents a single simulation run for 100,000 generations on a 300 × 300 grid. The final distributions of toxin producers for each species has been classified manually into a number of categories represented by different colours, according to the scheme described in section 9.4.1. Distributions of toxin production are given for the points marked with white circles in figure 9.7.

individuals of each species tend to cluster. Examples of distributions from each of the major groups are shown in figure 9.7. Each histogram in figure 9.7 shows the distribution of toxin-producers at one of the points marked with a white circle in figure 9.6.

The classification of the distribution into types is based only on whether individuals cluster at v = 0, v = 1, or at some intermediate level, so with two species and three possible clustering points the theoretical number of different types of distribution is $2^{2\times3} = 64$. Sixteen of those involve an extinction of one species, and of the remaining 48, only seventeen are observed in simula-



Figure 9.7: The coloured histograms show the trait (v_i) distributions for species *A* and *B* for the simulations corresponding to the eleven points shown by circles in figure 9.6. In each case the colour of the histogram matches the colour of the 'class' of distribution, as classified according to the scheme in section 9.4.1. The arrows \rightarrow



underneath the histograms show the most common transitions which occur in each of the eleven main outcomes, and were generated in the same manner as the similar diagrams in figure 6.8 (page 105).

tions. Eleven of the seventeen appeared in large areas of the parameter space and it is these eleven that are shown in figure 9.7.

9.4.2 Saturation of toxin production for one species

The theoretical model of equations (9.3) predicts that away from the critical value of γ , one of the two species will drawn towards all-or-nothing toxin-production, and this is the case in the distributions shown in the figure.

When γ is less than the critical value, the distribution of toxin-producers settles into a state like those shown in figures 9.7(a),(f)-(k). In these cases all of the species *B* individuals turn to nearly maximum toxin-production, in accordance with equations (9.3) which predict that \bar{v}_B should rise to one when $\gamma < \frac{1}{2}$.

When γ is larger than the critical value, the distribution of toxin-producers settles into a state like those shown in figures 9.7(c)-(e). In these three cases the stronger species settles into a state of zero toxin-production, in accordance with the prediction from equations (9.3) in which \bar{v}_A falls to zero when $\gamma < \frac{1}{2}$.

9.4.3 Evolutionary branching

When evolutionary branching and polymorphism occurs in the model, it usually happens in the species which does not have its mean toxin-production level saturated at zero or one. I will call the species whose trait is stuck at zero or one the restricted species. From figure 9.7 it can be verified that in every case the individuals belonging to the restricted species bunch up very closely at the extreme.

In some areas of the parameter space, the unrestricted species faces disruptive selection and splits into two (or sometimes even three) morphs, while in other regions, the unrestricted species stays monomorphic, with one cluster centred around a single value of v. Sometimes this mean value is intermediate between zero and one.

Monomorphic behaviour is shown in figures 9.7(a),(c),(e),(f),(k), and polymorphic behaviour is shown in figures 9.7(d),(g)-(j).
Monomorphic and polymorphic behaviour at high β

When $\gamma = 0.33$ and $\beta = 0.8$ (the light purple region in figure 9.1, figure 9.7(k)), *B* is the restricted species and all the *B* individuals are bunched up at v = 1, but the *A*s remain monomorphic with a concentration centred at $v \approx 2/3$.

As do not become dimorphic here because (1) fully toxin-producing As invade the Bs too slowly relative to the rate at which they are being invaded by the As with $v \approx 2/3$, and (2) non-toxin producing As invade the As with $v \approx 2/3$ too slowly relative to the rate at which they are invaded by the Bs.

If γ is increased, or β is decreased from these levels, then the *A*s become dimorphic. For example, when $\gamma = 0.43$ and $\beta = 0.8$, or when $\gamma = 0.33$ and $\beta = 0.6$ the *A* population diverges into two morphs; both of these cases fall into the grey region in figure 9.1, with distributions similar to the one in figure 9.7(h).

The reason for this is that an increase in γ means the toxin becomes weaker, so the *Bs*, which are restricted at v = 1, invade the *As* more slowly. *Bs* also invade more slowly when β is reduced: in this case both their toxin and their growth rate are weakened.

When these *Bs* invade more slowly, fully toxin-producing *As* are able to invade the *Bs* at a rate that is fast enough so that they can 'escape' from non toxin producing *As*, and non-toxin-producing *As* are able to invade toxin-producing *As* at a rate that is fast enough to escape from the *Bs*.

This boundary between the region where species *A* is polymorphic and the region where it is monomorphic appears to be robust, at least when $\gamma > 0.25$. It remains in the same location for values of σ from 0.0003 to 0.02, and for population sizes from 200 × 200 up to 1000 × 1000. This suggests that the monomorphic behaviour is caused by a lack of 'room' in the trait space around the optimal \bar{v}_A rather than by a lack of available diversity.

Monomorphic behaviour at very low β

The unrestricted species also becomes monomorphic when β is close to γ . This is seen in the dark red and yellow regions of figure 9.1, and figure 9.7(e)-(f).

In these cases the weak species, when producing its maximum amount of toxin, is a reasonably even match for the non-toxin-producing strong species. A toxin-producing morph of species *A* is not viable in this situation; its introduction would produce a RPS game in which it would be too vulnerable



Figure 9.8: A very simplified picture of the behaviour of the simple toxin-growth model as β and γ vary. The β - γ space is divided into three main regions, one in which the strong species always wins ($\rho \rightarrow 1$), one in which all weak species individuals produce a maximum amount of toxin ($\bar{v}_B \rightarrow 1$), and one in which all strong species individuals produce a minimum amount of toxin ($\bar{v}_A \rightarrow 0$). Within the latter two regions are subregions in which the species that has not fixated to zero or one is monomorphic (grey area) or polymorphic (white area).

because in RPS, its density can only be proportional to the rate at which the *B*s invade the non-toxin-producing *A*s, and this rate is very close to zero (in fact it is $\beta - \gamma$, as shown in figure 9.2).

Areas of cyclic competition

The arrows in figure 9.7 show that whenever at least one species becomes polymorphic, there is a very clear competitive cycle of some sort. Cyclic competition is the outcome in all cases in which the unrestricted species goes polymorphic, and also in the small region where γ is at its critical value and β is close to one (shown in figure 9.1 (the purple area) and figure 9.7(b)).

9.4.4 Summary of equilibrium behaviour

For the simplest toxin-growth model in which assumptions (9.1) and (9.2) hold, the main types of behaviour are summarised in figure 9.8. The main findings are:

(1) The two species can only coexist when the toxin is strong enough for a toxin-producer of the weaker species to outcompete a non-toxin-producer of the stronger species (white and grey areas in figure 9.8).



Figure 9.9: Species densities in simulations in which a single toxin-producing species *B* individual is introduced into a non-toxin-producing species *A* monoculture. Each point represents a simulation for up to 10,000 generations on a 300×300 grid. Extinctions are shown by dark brown and dark blue areas which indicate the winning species (*A* or *B* respectively). Coexistence after 10,000 generations is shown by all other colours, where the shade indicates the final species densities.

- (2) When the strength of the toxin is above some critical level, all weaker species individuals will produce toxin ($\bar{v}_B \rightarrow 1$ region), and strong species individuals will sometimes cluster into one group (grey subregion) and sometimes organise into more than one group (white subregion).
- (3) When the strength of the toxin is below the critical level, no strong species individuals produce toxin ($\bar{v}_A \rightarrow 0$ region), and weak species individuals will sometimes organise into one group (grey subregion) and sometimes organise into more than one group (white subregion).

Cyclic competition can evolve from the two-species system but only when the parameters are such that the system falls into one of the white areas in figure 9.8.

9.5 Invasion of an established population

Here I consider the situation in which a lone toxin-producing individual of the weaker species meets an established non-toxin-producing population of

172 CHAPTER 9. CYCLIC COMPETITION AMONG TOXIN-PRODUCERS

the stronger species. Figure 9.9 gives some idea of the likelihood of stable coexistence emerging out of such a scenario when $\gamma = 1/3$ and $\gamma = 2/3$.

Like the intraspecific competition model, there are really three types of extinction:

- (1) Immediate extinctions in which species *B* is unlucky and is wiped out at an early stage of the simulation before it manages to spread. These are indicated by the scattered brown dots which appear all over the graphs. The introduced toxin-producer is more likely to go extinct when β is low, because it will spread more slowly. It is also more likely to go extinct when γ is high, because its toxin is weaker.
- (2) Secondary extinctions occur after the species *B* individual multiplies early on, but species *A* fails to adapt by producing more toxin. These are indicated by the blue regions at the left of the diagrams. They are more likely when β is high (the invader is stronger), when σ is low (the resident adapts more slowly), and when γ is low (the invader's toxin is stronger).
- (3) A final round of extinctions occurs when species *B* is so weak that the system cannot reach an equilibrium involving coexistence, indicated by the solid dark brown areas at the lower right of the diagrams. These appear to be more likely when adaptation is quick, presumably because any move by species *B* away from full toxin production is detrimental.

An increase in the size of mutations reduces the likelihood of secondary extinctions by increasing the adaptability of the resident species, but it also increases the chance of the equilibrium extinctions by reducing the average toxin production of the weaker species.

9.6 Generalised toxin-growth model

The two toxins are unlikely to be exactly equal in strength, so in this section I will examine the more general model which relaxes the assumption in equation (9.2):

$$g_0/h_0 = g_1/h_1$$

Without this assumption, there is no longer a γ ratio, but there will still be some restrictions on the four parameter values g_0, g_1, h_0, h_1 . Firstly, the very notion of toxins implies that they are more dangerous than normal competition (or overgrowth), so I will assume that $g_0, g_1 \le h_0, h_1$. Secondly, I will call the species with the higher absolute toxin strength species A, and this allows the analysis to be restricted to the cases in which $h_B \le h_A$. So the only constraints of the parameters are

$$g_A \le h_B \le h_A, \qquad g_B \le h_B \le h_A. \tag{9.7}$$

It is now possible for the slower-growing species to have the strongest toxin $(g_A < g_B)$, which was impossible in the simpler model. If g_A and g_B are interpreted as the species' inherent competitiveness, or inherent efficiency in using the available resources (rather than as their 'growth rates') then the situation in which $g_A < g_B$ can be seen as a case in which an inherently weaker species can synthesise a chemical cheaply, or in relatively small amounts, but which happens to be particularly toxic against the stronger species.

9.6.1 Mean field equations

The mean field version of the general model is complicated by the fact that there are four parameters g_0, g_1, h_0, h_1 . It can described by the following equations:

$$\begin{aligned} \frac{d\rho}{dt} &= \rho(1-\rho) \left[g_A (1-\bar{v}_A) + h_A \bar{v}_A - g_B (1-\bar{v}_B) - h_B \bar{v}_B \right] \\ \frac{d\bar{v}_A}{dt} &= \mu((1-\rho)h_A - g_A) \\ \frac{d\bar{v}_B}{dt} &= \mu(\rho h_B - g_B) \end{aligned}$$

Essentially a species produces more toxin if its growth rate is less than its toxin strength multiplied by the other species' density: \bar{v}_A rises if $g_A < (1 - \rho)h_A$ and \bar{v}_B rises if $g_B < \rho h_B$.

The extra parameter in this system makes it difficult to see immediately which of the two species will become restricted when its mean toxin-production level saturates at zero or one. This information can be obtained from the simulation results presented in the remainder of this section.

9.6.2 Species densities in the general model

Figure 9.10 shows the equilibrium species densities over a range of g_A , g_B , h_B parameter values when h_A is fixed at 1. The species with the stronger toxin dominates when one toxin is more than twice the strength of the other (h_B <



Figure 9.10: Equilibrium species densities for the general model. The toxicity of the stronger species, h_A is assumed to be equal to 1, and each of the nine pictures shows results for nine values of h_B . Each point represents the species densities for one simulation after 10,000 generations on a 300 × 300 grid, starting with equal species densities and uniform levels of toxin production across individuals. The grey areas are combinations of the parameters that fall outside the conditions in the assumption (9.7), and the colours represent the densities of the species *A* according to the colour scale shown. The black lines show the plane in which the constraints of the simpler model of section 9.2 hold.



Figure 9.11: The small region of β - γ space in the simple model of section 9.2 in which the species with the weaker toxin can predominate.

 $\frac{1}{2}$). When $h_B > \frac{1}{2}$, the toxins are more equally matched and either species can wind up with a higher density.

When both toxins strengths are equal ($h_B = 1$), the densities depend on the growth rates. For most growth rates, the faster-growing species is at a disadvantage, because the production of toxin leaves it more vulnerable to its fast-growing conspecifics. But when both toxins are weak ($g_A, g_B > \frac{1}{2}$), then there is less incentive to produce it, and the faster-growing species has a higher density.

When one toxin is stronger than the other ($h_B < 1$), the situation changes and the species with the stronger toxin always has the highest density as long as the toxin is sufficiently strong ($g_A < \frac{h_A}{2}$).

Comparison with the simpler model

In the earlier model of section 9.2, the ratios of growth to toxin strength were fixed at the same value for both species. This constraint is shown by the black lines in figure 9.10, which show why the stronger species nearly always finished with the highest density in that model.

The blue areas in figure 9.10 occur under two conditions:

(1) In the upper right of the pictures, when $\frac{1}{2}h_A < g_A < g_B < h_B < h_A$. This could not occur in the simple model because when h_B is less than h_A , that model's assumption that $g_A/h_A = g_B/h_B$ necessarily implies that g_B is less than g_A .



Figure 9.12: Regions of g_A - g_B -space with qualitatively distinct outcomes for the generalised toxin-growth model when $h_B = 0.89$, $\sigma = 0.003$. Each square represents a single simulation run for 100,000 generations on a 300 × 300 grid. The final distributions of toxin producers for each species has been classified manually into a number of categories represented by different colours, and described in section 9.6.3. The distributions at the points marked with white circles are shown in detail in figure 9.13.

(2) In the upper left, when $g_B < \frac{1}{2}h_A < g_A < h_B < h_A$. This is possible in the simple model but only in the relatively small region of β - γ space shown in figure 9.11.

Although part of this second region is included in the simple model, in practice the behaviour in which the weaker species predominates is not noticeable in figure 9.1(a), probably because the densities are so close to $\frac{1}{2}$ in that region.

9.6.3 Distribution of toxin-producers in the general model

The species densities shown in figure 9.10 give the impression that the system behaviour varies relatively smoothly above and below the line where $g_A = g_B$. In fact the distributions of toxin-production can change dramatically with small variations in g_A , g_B even when the equilibrium densities (and mean trait

levels) are similar.

Regions with similar distributions are shown in figure 9.12 for the case shown in the bottom middle of figure 9.10 (where $h_A = 1$ and $h_B = 0.89$). Distributions of toxin-producers are classified into categories based on whether or not there is a dimorphism in either species due to disruptive selection. The dark brown stripes across the top and bottom, and the dark blue region in the upper left, are the regions with extinctions in which species *A* and *B* (respectively) win out completely.

Regions marked with other colours have qualitatively different distributions of toxin-producers. Examples of the most common types are shown in 9.13.

There are three main types of behaviour, each of which can be divided into a monomorphic and a dimorphic region. The diagonal going from upper left to lower right (separating the orange and red regions from the green and blue) is the generalisation of what I called the critical γ threshold in the simpler model. Below and to the left of this line, in the red and orange regions, species *B*'s toxin-production saturates at $\bar{v} = 1$, and species *A*'s mean toxinproduction level settles at some lower value. (Because species *A*'s toxin is always stronger, there is never an incentive for *A*'s toxin production to rise to one).

The diagonal 'critical γ line' appears from the simulations to occur at

$$\frac{g_A}{h_A} + \frac{g_B}{h_B} = 1,$$

and only determines whether or not \bar{v}_B saturates at 1.

Above the line, the toxin-production of one species falls to a minimum at $\bar{v} = 0$. Which species becomes restricted depends only on the relative values of g_A and g_B . If A grows faster than B (the two green regions), then A saturates at zero toxin-production and B's mean toxin-production level settles at an intermediate value; if B grows faster than A (the two blue regions), then B saturates at zero toxin-production while A settles at an intermediate value.

Looking at figure 9.10, we can see that within these regions, the density of the restricted species settles at a level that is equal to the growth rate of the unrestricted species. Figure 9.14 provides a summary the approximate behaviour of the system when extinctions are ignored.



Figure 9.13: The coloured histograms show the trait (v_i) distributions for species A and B for the simulations corresponding to the six points shown by circles in figure 9.12. In each case the colour of the histogram matches the colour of the 'class' of final outcome, as classified according to the scheme described in section 9.6.3. Arrows beneath the histograms show the common transitions using the procedure described in figure 6.8 (page 105).



Figure 9.14: A simplified picture of the three main regions of coexistence in the generalised toxin-growth model. One species' average toxin-production reaches its maximum or minimum, and the densities are then determined by the growth rate of the other species.

Cyclic competition

The three different saturation regions can each be divided into two parts, those in which the unrestricted species splits, and those in which the unrestricted species remains monomorphic. For the reasons discussed earlier in section 9.4.3, these monomorphic regions are the ones in which the mean trait values are close to zero or one, and there is not enough 'room' in the trait space for viable individuals to spread out into groups.

All the red, dark green and medium blue regions of figure 9.12 involve RPS ecosystems in which one species splits into a fully toxin-producing and a non-toxin-producing morph, as shown by the most common site transitions for these regions in figure 9.13(b), (d), and (e).

9.6.4 Invasion of an established population

In the more realistic scenario of an established, intraspecifically competing (non-toxin-producing) species *A* population being invaded by a lone toxin-producing species *B* individual, outcomes involving coexistence are also possible, but less likely than in the case when the simulations begin with equal species densities.

Figure 9.15 shows the results of simulations for 125 different sets of h_A , h_B ,



Figure 9.15: The final species densities for 300×300 grid simulations in which a lone toxin-producing species *B* individual is introduced into a resident population of non-toxin-producing species *As*. The grey scale on the left shows the values of the four parameters h_A , h_B , g_A , g_B , and each horizontal band on the right shows the final species densities after 10,000 generations for simulations which combine these parameter values with values of σ ranging from 0.0001 to 0.08. In the h_B scale, the five shades of grey represent fractions $(\frac{3}{8}, \frac{1}{2}, \frac{5}{8}, \frac{3}{4}, \frac{7}{8})$ of h_A , going from dark to light grey. The g_A and g_B scales represent fractions $(\frac{1}{6}, \frac{1}{3}, \frac{1}{2}, \frac{2}{3}, \frac{5}{6})$ of h_B , going from dark to light. The species densities use the same scale as those in figure 9.9 (page 171).

 g_A , g_B parameter values and for a range of 100 values of σ . Changing the mutation rate has little effect on the prevalence of extinction, but a large effect on which species is more likely to go extinct. Because *B* has the initial advantage, increasing adaptability primarily benefits the resident *A*.

9.7 Summary

The toxin-growth model assumes that

- (1) Individuals of two species are able to produce interspecific toxins in varying degrees;
- (2) All individuals are susceptible to the toxin produced by the other species; and
- (3) Individuals face a linear trade-off between their rate of toxin production and their growth rate.

The model shows that for a simple two-species system, when degrees of interspecific toxin-production are evolvable, variations in toxin strength and natural growth rates produce an incredibly diverse range of polymorphic and monomorphic equilibria.

Most small changes in the parameters lead to similar equilibria, because large areas in which the equilibria have the same properties are delimited by a few critical thresholds. The most important of these thresholds determine whether one or other species maximises or minimises its toxin-production. Within the regions defined by these thresholds, secondary parameter thresholds decide whether the other species becomes polymorphic or monomorphic. When a polymorphism occurs, this always produces intransitive competition with dynamic spatial effects.

Earlier models of toxin-based competition have emphasised the importance of spatial structure (Frank, 1994; Durrett and Levin, 1997; Iwasa et al., 1998), provided evidence for the maintenance of stable polymorphisms (Frank, 1994), intransitive three-species cycles (Durrett and Levin, 1997), and complex intransitive webs (Pagie and Hogeweg, 1999; Czárán et al., 2002). The model described in this chapter suggests a mechanism for the origin of polymorphisms, which may help to explain the incredible diversity of ecosystems of toxin-producers.

182 CHAPTER 9. CYCLIC COMPETITION AMONG TOXIN-PRODUCERS

Part IV

Competitive restraint

Chapter 10

Community level selection as an explanation of competitive restraint in RPS ecosystems

One interesting aspect of RPS ecosystems is the evolution of competitive restraint discovered by Johnson and Seinen (2002). In an experiment which takes place on a SCA model of a RPS ecosystem, the invasion rate of one of the three species (rock) is allowed to evolve, and the invasion rates of other two species, scissors and paper, are fixed at their initial levels. As the grid is updated, the mean invasion rate of rocks increases up to a point but eventually stops rising. Although there is selection pressure for rocks to increase their speed of invasion in order to compete against other rocks, when they become too competitive, they deplete all the scissors individuals in their local area of the grid and are slowly overrun by paper. So the competitiveness of rocks eventually becomes constrained.

The competitive restraint that evolves in this model is an example of individual altruism, at least in the short term. Rocks restrain their competitiveness, thereby producing fewer offspring than they otherwise would have, and the primary beneficiaries of the restraint are completely unrelated individuals of a different species (scissors).

Johnson and Seinen claim that the evolution of restraint in these ecosystems is a case of *community level selection*. Natural selection between individuals towards greater invasiveness is competing with a second process of natural selection at the level of the ecological multispecies subcommunity, which selects among subcommunities for the ones that are the most stable. In section 10.1 I explain the theory of community level selection and how it has arisen from group selection and multilevel selection theory. Section 10.2 outlines some of the difficulties involved in defining higher level units of selection. Sections 10.3 and 10.4 describe aspects of the process of community level selection on spatially continuous communities and cellular automata. Finally in section 10.5 I argue that the invocation of community level selection to explain the phenomenon of restraint is an unnecessary overcomplication of the process, and that a better understanding of the phenomenon can be achieved by appealing to the details of RPS interactions and the spatial structure of the system, without the need to propose competition between subcommunities.

10.1 Levels of selection

There has been plenty of debate in biology regarding the levels at which natural selection acts, or the 'units' of selection. Much of the argument has centred around whether selection on genes or selection on organisms should be invoked in order to explain various adaptations, and these issues are not important here.

10.1.1 Group selection

The extent to which selection acts upon units bigger than individual organisms is more controversial. Traditionally this debate has been around group selection rather than the community level selection process claimed by Johnson and Seinen, but the processes are very similar.

Group selection is the idea that groups of individual organisms of the same species, such as wolf packs or ant colonies, have different levels of fitness, and that natural selection operates on these groups 'directly', in addition to operating on the individuals which make them up. Group selection has been regarded until recently as a discredited theory in evolutionary biology, primarily due to a controversial formulation of the theory by Wynne-Edwards (1962), which proposed that individual animals could evolve birth-control mechanisms that decrease the size of the individual's group and therefore prevent over-exploitation of local resources. These mechanisms were said to spread because of the resulting benefit to the group. Subsequent critiques of Wynne-Edwards' formulation by Maynard Smith (1964), Williams (1966) and

Dawkins (1976) then led to an overreaction against the theory of group selection in general (Wilson and Sober, 1994).

10.1.2 Group selection and altruism

Group selection has mainly been invoked to explain the existence of altruistic traits. For example, consider a gene that codes for altruistic behaviour. In other words, the presence of the gene decreases the owner's fitness and simultaneously raises the fitness of the owner's group, relative to groups which contain only selfish individuals. Such a gene can theoretically spread and become stable when the population has a structure in which altruists tend to associate together except during mating (Sober and Wilson, 1998). This is possible because even though selfish individuals will do better within each of the groups, the groups which contain altruists will have many more members than the selfish groups by the time the entire population comes together for mating.

10.1.3 Multilevel selection theory

The levels of selection debate has largely been resolved, and it is now commonly accepted that natural selection can and does operate on a nested hierarchy of units such as groups, organisms, and genes (Wilson and Sober, 1994). Differential fitness between groups can have a strong effect on evolution, and in simple models the contributions of group selection and individual selection to evolution can be separated and quantified (Price, 1970; Hamilton, 1975).

However, there is still debate over the most appropriate explanations for particular adaptations. There is a conflict between rival explanations because adaptations due to selection acting on a higher-level unit can also be explained solely in terms of selection acting on lower-level units.

Even proponents of multilevel selection theory generally agree that while all instances of selection can be said to act at the gene level, only a subset act at the level of the organism, and only a subset of those will act at the level of the group. But they also claim that there are instances of adaptations for which selection at the higher level should be the preferred explanation because such an explanation provides a deeper understanding of the processes involved (Sober and Wilson, 1998).

10.1.4 Group selection or individual selection?

A group selection interpretation of the altruistic adaptation in the hypothetical population described in 10.1.1 is that the altruistic gene spreads because it benefits the group, and groups of altruists do better than selfish groups in the struggle between groups. The same process can be explained without recourse to group selection by packing the population structure up and describing it as a part of the individual's environment.

The different contributions that altruistic and selfish groups make to the population may mean that average altruist has a higher fitness than the average non-altruist, when averaged across all groups. With such an explanation there is no need to describe individual behaviour as 'altruistic' because individuals are now maximising their self-interest within the context of a population structure that happens to involve groups. Sober and Wilson would say that this interpretation is rather contrived, and that it is only favoured by those who for some reason cannot bear to admit to the existence of altruism. They would claim that a group selection interpretation of the process leads to a deeper understanding of it.

10.1.5 From group selection to community selection

Community level selection works in exactly the same way as group selection except that members of the 'groups' can be of different species. This means community selection is theoretically well understood. But although selection at the community level has been demonstrated in the lab using artificial selection (Goodnight, 1990a,b), according to Wilson and Swenson (2003), it has not yet been seen operating in nature.

Community level selection is usually described as a process whereby multispecies communities form, compete with one another, and sometimes are said to give birth to new communities which inherit some of the same community traits as their parents. An individual's altruistic trait, which benefits that individual's community at a direct cost to its own fitness, can spread through the population because its community gives rise to more 'child communities' in comparison to communities full of non-altruists.

10.2 Population structure

For selection to occur at higher levels, the overall population needs to be structured: it must include groups or subcommunities. A perfectly well-mixed population of individuals interacting at random will not support the evolution of altruism under most definitions.

10.2.1 What is a group?

There are different opinions on exactly how well-structured into groups a population must be before group selection can operate. In early group selection models, it was thought that groups needed to be relatively discrete, and stay together for a significant period of time (Williams and Williams, 1957), or even for several generations (Maynard Smith, 1964).

On the other hand, Sober and Wilson (1998) suggest that group selection could operate on groups whose members are not in spatial proximity to one another (if group members can recognise one another), or on groups with fuzzy boundaries caused by limited dispersal.

Groups, according to Sober and Wilson, should be defined on the basis of interactions, and only with respect to a particular trait. So for each individual trait, the set of individuals which will benefit from that trait, given the nature of individual-individual interactions, can be seen as the group for that trait. Even ephemeral associations of size two could be valid groups under this definition.

10.2.2 What is a (sub)community?

A similar set of opinions is evident among proponents and detractors of community level selection.

Maynard Smith and Szathmáry (1995) suggest that ecosystems cannot be units of selection because they lack the individuality and separateness found in organisms and social groups. But as they themselves point out, it is very widely believed that existing complex life forms such as multicellular organisms have in fact evolved out of multispecies communities into well-integrated assemblages. In other words, individuality and separateness can evolve. What is under dispute is how separate a community must be for it to count as a unit of selection.

190 CHAPTER 10. RESTRAINT AND COMMUNITY LEVEL SELECTION

In a recent set of articles about 'community genetics', Collins (2003) describes the views of community selectionists as lying on a continuum, similar to the range of views of group selectionists. At one end, there are those like Whitham et al. (2003), who think the likely candidates for community level selection are long-term, tightly-knit associations of species with well-defined community level traits. Most provisional examples of community level selection given in the literature involve these types of communities.

At the other extreme are those like Neuhauser et al. (2003), who believe that communities in only brief association can be acted on by selection. Johnson and Seinen fall closer to the latter view, because the competing subcommunities in their model, which they describe as undergoing community level selection, are neither discrete nor particularly long-lasting.

There is a trade-off here, because while tightly-knit communities are more likely to have the properties that will make community level selection a significant force, there are fewer of these tightly-knit communities among the ecosystems in nature. And for those who see nearly any loose association is a potential community, in most of these potential communities, we can expect that it will be difficult to find between-community selection which is sufficiently strong to counter the individual level selection forces.

10.3 Spatially continuous landscapes

Subcommunities within spatially continuous communities, in which limited dispersal provides the only population structure, are perhaps the most contentious candidates for community-level selection. If subcommunities have some degree of discreteness and coherence, then predatory restraint can evolve. This has been shown in previous simulations using metapopulations (Gilpin, 1975a; Mitteldorf et al., 2002). But in a spatially continuous landscape with limited dispersal, there is less coherence, because no two individuals have the same subcommunity.

Johnson and Boerlijst (2002), and Wilson (1997) believe that the limited dispersal provides enough separateness for the subcommunities to be considered units of selection, but Maynard Smith and Szathmáry (1995) clearly do not. This is despite the fact that altruism has been shown to evolve under these conditions.



Figure 10.1: Evolution of the invasion rate of rock (r_r) in two spatial rock-paperscissors simulations in which the invasion rates of scissors and paper (r_s, r_p) are held fixed at 0.2. In the first simulation (black line), r_r starts with a mean of 0.2 (ranging from 0.1-0.3), and in the second simulation (blue line), r_r starts with a mean of 0.7 (ranging from 0.5-0.9). In both cases r_r evolves to around 0.6.

10.3.1 The evolution of altruism in viscous populations

Altruistic traits are known to evolve on spatially continuous landscapes, with limited dispersal but without discrete group or community structure. Hamilton (1964) claimed that limited dispersal can promote altruism through kin selection: the intuitively appealing explanation is that limited dispersal increases the proximity of relatives, and because individual-individual interactions are more likely to be between relatives, genes for altruistic behaviour are likely to benefit copies of those genes in neighbours.

However, in models of fixed size populations with limited dispersal (called 'viscous' populations in most of the literature), it has been shown that limited dispersal encourages within-neighbourhood interactions, but also limits between-neighbourhood interactions to such an extent that the overall effect on the spread of altruism is neutral (Queller, 1992; Taylor, 1992; Wilson et al., 1992).

10.3.2 Fluctuations in population size

If the overall population size is allowed to fluctuate, however, the situation can be different — altruistic traits can spread, even when limited dispersal provides the only spatial structure (Mitteldorf and Wilson, 2000).

But the predatory restraint in RPS systems is quite different. It is not the sort of altruism that requires a fluctuating population in order to evolve. The Johnson and Seinen experiments which show competitive restraint used a fluctuating population, because their SCA grid allowed empty cells, but a variable population is not required. Competitive restraint can also be shown to evolve on similar experiments in a fixed population without empty cells. The results of two such experiments are given in figure 10.1: the invasion rate of the evolving species shows the same essential behaviour as that in figure 2a of Johnson and Seinen (2002).

10.3.3 Altruism and competitive restraint

If community level selection is responsible for the spread of competitive restraint in fixed size RPS communities, then there must be some reason why restraint is not subject to the same cancellation effect of within- and betweencommunity interactions that occurs for the altruism traits in other models such as Queller (1992) or Mitteldorf and Wilson (2000).

I believe that the nature of competition in the three-species intransitive network is the reason. The 'survival of the weakest' phenomenon in the RPS ecosystem usually guarantees that faster invaders are disadvantaged in all but the very shortest of timespans. Therefore restraint is not really altruistic in the usual sense, given the other competitors in the system.

This important difference between competitive restraint in RPS and group selected altruism becomes especially clear when we consider the alternative strategies to altruism and restraint respectively. In the altruism model, selfish individuals thrive in the absence of any population structure. The same cannot be said for 'unrestrained' competitors in RPS: with global interactions, if the invasion rate of rock, say, is allowed to evolve, then it will rise without limit until all the scissors are gone, after which rock soon goes extinct. Selfishness does well where unrestrained competition does poorly.

10.3.4 Provisional examples of selection in continuous communities

The most convincing descriptions of community level selection involve quite discrete-bordered communities in which the members of the community have a shared interest. There are numerous examples of this in the literature, including the view of single organisms as multispecies communities, symbioses, and species that use another species for transport between areas of favourable habitat (Wilson, 1997).

But there are fewer convincing provisional examples of spatially continuous

community selection. Wilson (1997) provides two, one involving combined plant-and-soil communities, and another suggested by Leigh (1994) in which root grafts between neighbouring trees of different species protect both trees against storms.

These examples are speculations and not intended to provide proof of the existence of community-level selection in nature, but in any case neither example shows quite the lack of community separateness in time and space that occurs in a purely spatially continuous landscape or in a cellular automata model. In the plant-and-soil example, the plant affects bacterial communities relatively evenly throughout a certain region of the soil under it, so there is some degree of long-range dispersal in the system. In the case of the root grafts, there may be long-range effects because we expect that neighbouring trees experience a storm equally. But the grafts themselves also have the effect of adding a degree of structure to the space which is not possible in a SCA, for example. Even if community selection turned out to be a good explanation for these provisional examples, that would not automatically show that the same effect could happen in a world with no variation in habitat.

10.4 Communities in cellular automata

In many ways the cellular automata grid used in the models of Johnson and Seinen (2002) is the ultimate spatially continuous landscape. Not only do subcommunities in these models lack separateness, but also the nature of interactions on the grid means there is no opportunity for a more complex population structure with separate subcommunities to evolve.

Subcommunities on a SCA grid must therefore be constructed in some way by the observer so as to have the right composition for community selection to act. Johnson and Boerlijst (2002) provide a list of properties that an ecosystem must have for community level selection to be a significant force:

- (1) Subcommunities must maintain a degree of individuality in time and space, and
- (2) There must be heritable variation among subcommunities.

They believe that the Johnson and Seinen (2002) model meets these criteria and that subcommunities on the SCA grid have a degree of separateness and individuality.

In this section I will describe how the community level selection process promotes competitive restraint, and explain why the lack of separateness of SCA subcommunities causes problems for a satisfactory explanation of competitive restraint.

10.4.1 The community selection process

The process that Johnson and Seinen describe involves a conflict between individual selection and subcommunity selection. In a RPS population, the invasion rate of rock is allowed to evolve while the invasion rates of scissors and paper are fixed. Selection at the individual, within-subcommunity, level favours an increase in rock's invasion rate without limit. If individual selection were the only force, rocks would invade faster and faster, eventually wiping out the scissors population after which the rocks themselves would quickly be eliminated by the remaining paper population.

However, at the same time, between-subcommunity selection drives the invasion rate of rocks down. The claim is that subcommunities compete on the basis of 'persistence stability', which is the property that all species in the subcommunity persist after one turnover of all individuals (Dayton et al., 1984; Johnson and Mann, 1988; Johnson and Boerlijst, 2002). Subcommunities with relatively even invasion rates between species have more persistence stability, and tend to outcompete, subcommunities with uneven invasion rates and less persistence stability.

Johnson and Seinen regard community level selection as a good explanation of the restraint phenomenon. However, it is a very limited explanation, because the relative strength of the two selection processes is unmeasurable without a reasonably well-defined concept of the subcommunity. Because there is no separateness, there is no obvious method for finding out the extent to which the fittest, most stable subcommunities produce more offspring than the less fit ones.

I now look at two potential ways of bringing clarity to the process by defining subcommunities on the SCA.

10.4.2 Community definition using interactions and traits

I have mentioned the fact that Sober and Wilson think groups should be defined on the basis of interactions, and with respect to a particular trait. This presents a special problem when translated to the case of defining SCA subcommunities because all interaction groups will overlap on a SCA grid.

In multilevel selection theory, a group is "a set of individuals that influence each other's fitness with respect to a certain trait, but not the fitness of those outside the group" (Sober and Wilson, 1998). In terms of community rather than group selection, if we take this definition to mean direct rather than indirect influence, there are two obvious ways to define subcommunities on a SCA, neither of which is entirely satisfactory:

(1) Ephemeral subcommunities of size two

In a SCA model with limited dispersal and asynchronous updates, only one pair of neighbours interact directly at any one time. The result of this the subcommunity equivalent of the "ephemeral groups of size N = 2" which are valid groups according to Sober and Wilson.

With each pair being a competing subcommunity, under the proposed model of community level selection the fittest pairs are those pairs with the most persistence stability, according to the Johnson and Boerlijst (2002) definition. But persistence stability is not meaningful for pairs. This is because with pairs in isolation, a complete turnover of all individuals is impossible in the RPS game – at least one of the individuals in the pair will be uninvadable.

(2) Grid neighbours as subcommunities

Johnson and Seinen's SCA uses synchronous updates throughout the entire grid, so it is more appropriate to consider a site and all its directly influenced neighbours as a subcommunity. With immediate neighbours it is possible to have all three species represented in the subcommunity, so the persistence stability of these subcommunities can be compared.

This can be done by running simulations on these tiny subcommunities until either an extinction occurs or all individuals have turned over at least once, and then counting the number of species left (provided periodic boundary conditions are added, or spatial effects are removed, from the subcommunity, because otherwise some individuals on the edge will be uninvadable). With enough samples the persistence stability of several sets of invasion rates and initial conditions can be quantified.

A disadvantage of such small subcommunities is the granularity in initial species densities. With four neighbours per site, the subcommunities will have a maximum of five individuals, and so some combinations of initial den-

196 CHAPTER 10. RESTRAINT AND COMMUNITY LEVEL SELECTION

sities, such as all species in equal proportions, are not possible.

Problems with small subcommunities

These conceptions of subcommunities conform to the Sober and Wilson definition, because all members of the subcommunity interact and affect one another's fitness. However they are limited in their ability to explain the restraint phenomenon for two reasons.

Firstly, there are at least as many possible overlapping subcommunities as there are sites on the grid. Between-subcommunity interaction in such a situation is extremely complex. Any attempt to reduce this complexity by restricting the analysis to non-overlapping subcommunities would require arbitrary decisions about which subcommunities were 'separate' which would fail to reflect the structure of the underlying model.

Secondly, with tiny subcommunities, it is extremely unlikely that any of them will show persistence of all three species for more than a handful of generations. The detail of figure 2.5 on page 21 shows that it is rare to find all three species present in very small areas. There is just too much same-species clustering, because at these scales individual selection dominates.

10.4.3 Community definition using spatial self-structuring

Johnson and Boerlijst (2002) believe that the individuality of subcommunities is emergent in the model through a process of spatial self-structuring, by which they mean the tendency of spatial RPS systems to organise into clusters, with a maximum size, which continually move across the grid. This structure could potentially be used to define the extent of the subcommunities.

The simulations of chapter 3 showed that stability continues to increase as the grid size is increased. For any particular time horizon, there will be some scale at which coexistence of all three species effectively becomes stable over that time period.

There will also be a grid size for which the probability of stability, for whatever timescale it takes for a single turnover of individuals, crosses some threshold, and this grid size could be used to define the size of the subcommunity. Such a threshold will necessarily be arbitrary, because stability continues to increase with every increase in subcommunity size.

These subcommunities will also have to be a different size for every different

set of fixed invasion rates. This fact makes community level selection next to useless for predicting how restrained a species in a particular ecosystem will become: simulations will be necessary to find out the size of the subcommunity, and if simulations are required, then the degree of restraint can be found directly from the simulations.

10.4.4 Subcommunity interaction and heritability

Because interactions between any two adjacent sections of a SCA grid are complicated, it is difficult if not impossible to describe the way in which subcommunities compete and reproduce in the simple manner by which it can be done for discrete-group models.

On a SCA, it is not obvious that individuals from more stable subcommunities will leave more offspring than those from the less stable subcommunities. For example, consider two subcommunities from the description of the restraint process given in section 10.4.1. In the first, rocks are restrained and invade only a little bit faster than the other two species, and in the second, rocks are unconstrained. When these two subcommunities interact, rocks from the second subcommunity will probably be wiped out or be vastly reduced in number, but paper individuals from the second subcommunities. The only way to find out for sure is to keep track of individual lineages (as opposed to subcommunity lineages) in the SCA, because the process of subcommunity inheritance is just too complicated.

10.5 Conclusion

Not even the most enthusiastic of group selectionists believe that higher level units should always be invoked when explaining a particular adaptation. According to Williams (1966), we should not refer to adaptations of the group unless a process of natural selection can be shown to act on those groups. Sober and Wilson agree, and in fact it was the assumption by Wynne-Edwards (1962) of group-level adaptations without a corresponding process of grouplevel selection that was partially responsible for giving group selection a bad name.

Similarly, community-level adaptations will only arise when there is a process of natural selection acting on communities, for which communities must act

198 CHAPTER 10. RESTRAINT AND COMMUNITY LEVEL SELECTION

as functional units to some degree (Wilson, 1997). The degree to which subcommunities in a spatially continuous community can be regarded as functional units is debatable, but in any case, an attempt to *show* that there is a process of between-subcommunity selection going on should at the very least try to identify what the subcommunities are. Neither Johnson and Seinen (2002) nor Johnson and Boerlijst (2002) offers a clear identification.

But the true test of whether or not multilevel selection theory should be invoked is whether or not the theory "provides an added understanding of the underlying mechanism" that selection at the lower level does not (Johnson and Boerlijst, 2002). The kind of underlying mechanism that is better understood by the application of multilevel selection theory is the existence of two "separate causal processes that contribute to the evolutionary outcome" (Sober and Wilson, 1998). Multilevel selection theory is especially useful when the processes are opposing.

In the RPS model of competitive restraint, there are opposing processes at work, but I believe that regarding those processes as individual level and community level selection does not add to our understanding of the system, because the subcommunities themselves are too vague and too poorly understood. One can achieve a better understanding of the phenomenon by describing the opposing processes as the short term and long term self-interest of the competitor showing restraint, given the structure of RPS interactions. In the simplest terms, those that invade too quickly will deplete the enemy of their enemy and be overrun, and those that invade too slowly will be beaten by their neighbours.

An account of the details of the 'survival of the weakest' property of RPS ecosystems is required here in order to explain why long-term and short-term interest are in opposition. But with a community selection description, the details of RPS need to be a part of the explanation anyway, because they are the only way to make sense of how invasion rates affect persistence stability.

Chapter 11

Competitive restraint and the measurement of stability

Johnson and Seinen (2002) claim that the phenomenon of competitive restraint in spatial RPS systems is caused by tension between two opposing selection processes: selection for increased competitiveness at the individual level versus selection for stability at the level of the ecological community. Their restraint experiment takes place on a SCA grid of competing individuals of three species in a RPS-like relationship, in which the invasion rate of one species is able to evolve but the invasion rates of the other two are fixed. Eventually, if all three species survive, the average invasion rate of the evolving species reaches an equilibrium that is less than its theoretical maximum, and less than what it could be if individual selection were the only force.

While the evolving species is said to 'restrain' its competitiveness at this equilibrium level, the selective forces are of course not acting directly on species: the process envisaged by J&S is one whereby individuals that are too competitive tend to die out, because they make the subcommunities to which they belong so unstable that those subcommunities collapse into a monoculture which fares poorly in comparison to other subcommunities.

If this picture is correct, when the evolving species is at its equilibrium competitiveness level there should be many subcommunities collapsing, and it should be possible to observe the resulting instability in the system. In particular we might expect the system's instability to reach some limit above which it cannot rise as it becomes moderated by the restraint of the evolving species.

In an attempt to find evidence for this phenomenon, I examine the stability of RPS ecosystems in three stages:

(1) Measurement of community stability

The competitiveness trait acted on by individual-level selection is already well-defined within the system. But the community-level trait that has been proposed as the object of the higher-level selective force, stability, can also be measured in the SCA model in various ways. I propose four different measures of community stability, which are examined in turn in the next four sections. These four measures have slightly different characteristics, but agree for the most part with the expectation by J&S that systems with relatively similar invasion rates are more stable than those with uneven invasion rates.

(2) Strength of individual-level selection

By calculating the stability of the system using a number of sets of fixed invasion rates and then comparing this to the stability of the system using the equilibrium sets of invasion rates from the restraint experiments, it is possible to build up a picture of just how far away the system settles from maximum community stability. The distance provides an indication of the strength of the individual-level selective force, if the simple two-process model is accurate.

In accordance with the J&S model, in most cases the rate of the evolving species is not far from halfway between its theoretical maximum (the endpoint of individual selection) and the level that would produce the most stability (the endpoint of community selection). The four measures of community stability usually settle between the endpoints of the two kinds of selection (maximum stability and maximum competitiveness), but the equilibrium stability levels are more complex and less predictable.

(3) Variation in equilibrium stability

Finally I compare the stability of the system at several sets of invasion rates corresponding to a range of different equilibrium restraint levels in the J&S experiments. If there were a simple mechanism by which subcommunities go extinct once their invasion rates become too uneven, as the J&S picture suggests, then a consistent level of instability might be expected at the equilibrium invasion rates. None of the four types of stability turns out to be constant across the equilibrium, so this basic analysis provides no confirmation of any simple model of two opposing processes of selection working at different levels.



Figure 11.1: A comparison of the longest average time to extinction with the equilibrium points of the restraint system described in Johnson and Seinen (2002). Each dotted line shows one ratio of r_s to r_p , and the shaded circles on the dotted lines show the TTE (averaged over 2500 simulations on a 30 × 30 grid) for the combination of invasion rates represented by the location of the circle. White circles indicate the longest TTEs, and black circles the shortest. The solid black line joins the points with the maximum average TTE for each ratio of r_s to r_p . The equilibrium invasion rates in the restraint system where r_r is allowed to evolve are shown in red. These were measured on a 500 × 500 grid with initial $r_r = 0.5$ and r_s , r_p fixed in the appropriate ratios such that $r_s + r_p = 0.5$. Each simulation was run for an initial 12,000 generations and the equilibrium r_r was found by taking the average for all rock individuals over a further 60,000 generations. All points on the diagram are normalised so that $r_r + r_s + r_p = 1$.

11.1 Time to extinction on small grids

In chapter 3 the stability of SCA ecosystems was measured using time to extinction (TTE), by counting the generations until one species takes over the entire grid. This measurement closely corresponds to the notion of persistence used by ecologists because it is concerned only with the presence of species and not their densities. While it is easy to measure time to extinction, the high variability between simulations means that it is necessary to average TTE over multiple simulations with small random differences in the starting conditions so that the measurement is useful for comparing stability at different sets of invasion rates.

A fairly small grid must be used when measuring TTE in order to ensure that an extinction happens in a reasonable time. But the use of small grids is not a problem for the assessment of the J&S model, because small grids are just like subcommunities of a large SCA community, so TTE on small grids corresponds to the notion of *subcommunity* persistence, the trait which is supposedly being selected for.

11.1.1 Invasion rates and time to extinction

Figure 11.1 shows how TTE varies with combinations of invasion rates. Each point on the triangle represents a set of invasion rates, and the system stability is represented by the shading in the circle, with stability increasing from black to grey to white. The average TTE on the small 30×30 grids was around 1500 generations, ranging from a minimum of around 100 generations for the rates (0.5, 0.425, 0.075) to a maximum of around 370,000 generations for the rates (0.35, 0.325, 0.325). The white in the centre and black at the edges show that the system is more stable when invasion rates are evenly matched.

Dependence of TTE on ordering of invasion rates

There is a smaller triangle, made up of grey points, visible inside the larger triangle of black points, and this smaller triangle is not exactly centrally spaced within the larger triangle; it appears rotated slightly to the left. The offset in this area of increased stability means that two systems with the same unordered set of uneven rates may be more or less stable depending on which species those rates happen to be assigned to. Those systems in which the slowest invader invades the fastest (for example when $r_r > r_s > r_p$) will survive longer than systems where the fastest invader invades the slowest (for example when $r_r > r_s > r_p$).

In the latter example, rocks are the most vulnerable species because their equilibrium density corresponds to the slowest rate r_s (see the predictions of the mean field equations on page 3). Because rocks are being predated upon faster here than they are in the former example, they tend to go extinct more easily. Additionally, in the former case, even rocks that have become surrounded by

202

paper can be 'rescued' by groups of scissors that invade paper faster than paper invades rock. This rescue effect is unlikely when scissors invade paper more slowly than paper invades rock.

11.1.2 Individual and subcommunity selection

Each red point in figure 11.1 shows the equilibrium invasion rates if r_r is allowed to evolve while r_s and r_p are fixed at the ratio marked by the dotted line that passes through the point. For each of these r_s : r_p ratios, the system could settle at any point along the dotted line, but is only stable (without extinction) at the red point.

The solid black line joins the points at which average TTEs are at their maximum for each of the $r_s:r_p$ ratios. Under the J&S model in which subcommunities are selected for stability, multilevel selection theory says that if community level selection were the only force, then the system would always settle somewhere close to this solid black line. The distance between the red line and the black line therefore represents the strength of individual-level selection, which is the only opposing force.

If individual-level, rather than community-level selection were the only force, then r_r would climb without limit and eventually overwhelm the other two invasion rates, so the distance between the red line and the point at the top of the triangle represents the strength of community level selection under the model.

Non-opposition of selective forces in extreme cases

At first glance the location of the red line in figure 11.1 shows no evidence of selection for stability. The dotted lines at the right hand side of the diagram represent ratios of $r_s:r_p$ that are somewhat unstable regardless of the invasion rate of rock, because all the circles on these dotted lines are quite dark. Despite this instability however, rock responds by becoming especially competitive, as evidenced by the upward kink in the red line on the right hand side. It appears that the pull towards stability is weak in a naturally unstable environment when selection between subcommunities ought to be particularly intense because large numbers of subcommunities will be collapsing to monoculture and being 'invaded' by more diverse subcommunities.

For these extreme $r_s:r_p$ ratios, however, the time to extinction behaves in an unexpected way. Figure 11.2 shows the stability surface in detail as a contour



Figure 11.2: Contour plot of time to extinction versus invasion rates using the same data as figure 11.1, where the red line shows the restraint equilibrium and the black line maximum stability. The *x*-axis shows ratios of r_s and r_p , the *y*-axis values of r_r . Black areas indicate early extinctions, and white areas high stability. When r_s and r_p are very uneven, extinction times increase as r_r increases from the equilibrium line.

plot. From this plot it can be seen that stability does not always decrease monotonically as r_r moves away from the point of maximum stability. The top left and top right of the diagram show stability increasing with increasing r_r .

The implication of this is that the two selective forces, towards increasing rate of invasion and increasing stability, do not in fact oppose one another for some invasion rates even when r_r is very high. Therefore, if the restraint of r_r at this equilibrium is caused by stability at all, it is caused by some kind of stability that is not contained within the TTE measurement.

TTE is the only one of the four stability measures considered in this chapter for which invasion speed and stability are not in opposition at the equilibrium.


Figure 11.3: TTE averaged over 2500 simulations on a 30×30 grid, for a range of $r_s:r_p$ ratios. The dotted line shows the maximum average TTEs, at the points joined by the solid black line in figure 11.1. The solid line shows the average TTE for simulations run with invasion rates at the equilibrium level (the red points in figure 11.1) when r_r is allowed to evolve.

11.1.3 Variation in extinction times at the restraint equilibrium

The actual values of TTE measurement along the red and black lines are shown in figure 11.3. TTEs at the invasion rates where r_r has reached equilibrium vary by around two orders of magnitude across the range of $r_s:r_p$ ratios. TTE experiments therefore provide no evidence that community level selection is pulling stability above some threshold of subcommunity survival.

11.2 Minimum species density

The second measure of community stability is the minimum density of the most vulnerable species over time, the same as the definition of extinction risk from section 5.2 on page 80. This measurement is expected to correlate with the persistence of the three-species ecosystem because in general smaller populations are at greater risk of going extinct than larger ones.

Figure 11.4 shows the extinction risk measured throughout the space of possible invasion rates. The extinction risk is low for evenly matched invasion rates and high for uneven rates. This is unsurprising because we know that in spatial RPS the three species end up occupying space in proportion to the three invasion rates.



Figure 11.4: Minimum extinction risk (black line) compared against extinction risk at restraint equilibrium (red line). The shading in a circle indicates the minimum density of the least common species during the final 12,000 generations of a 24,000 generation RPS simulation on a 500 × 500 grid for a fixed set of invasion rates. A white circle means the extinction risk is small because the least common species has a density close to 1/3; a black circle means the extinction risk is large, and the density of least common species is close to zero. The solid black line joins the points where the minimum species density is at a maximum for a particular ratio of r_s to r_p . The equilibrium invasion rates in the restraint system where r_r is allowed to evolve are shown in red.

The measurement is taken by running a simulation for a very long time and returning the lowest species density recorded during that time. This will be a number between zero and 1/3, the maximum possible density for the least common species.

The amount of time that the simulation spends near the minimum density is ignored. So for example, any two simulations in which rocks fall to a minimum of 1% of the population for a single generation will always have identical scores, even if one simulation spends only a few generations with rocks under 2%, and the other spends thousands of generations with rocks below



Figure 11.5: Extinction risk measured by the minimum density of the rarest species over time described in figure 11.4, for 15 ratios of $r_s:r_p$. The minimum densities for the equilibrium restraint invasion rates are shown by the solid line, and the maximum minimum density (lowest extinction risk) is shown by the dotted line.

2%. Despite the apparent crudeness of this notion of stability, it captures the 'distance' from extinction over time better than a measurement based on the mean species densities because extinctions only happen once and they're permanent.

The minimum density statistic does not have the property of the TTE (noted in section 11.1.1) in which stability is dependent on the ordering of the three invasion rates. Minimum density appears to depend on the unordered set of invasion rates but is relatively insensitive to whether the fastest invader preys on, or predates on, the slowest. So while 'fast eats slow' ecosystems tend to go extinct more quickly than 'slow eats fast' ecosystems, their minimum densities in figure 11.4 are similar, because they are measured with a much larger grid on which variations in the numbers of local extinctions are less noticeable.

Figure 11.5 gives the actual values of the measure at the maximum and equilibrium. It shows that there's no characteristic extinction risk at the equilibrium set of rates; the equilibrium stability follows the general shape of the maximum.



Figure 11.6: Comparison of the invasion rates which produce minimum density variation over time (black line) with the equilibrium invasion rates when r_r evolves freely (red line). The shading in the circles represents the average, over all three species, of the amount by which their densities fluctuate over time. This is measured by taking samples of species densities from the final 60,000 generations of a 72,000 generation simulation on a 500 × 500 grid, and averaging the coefficient of variance (the ratio of standard deviation to the mean density over time) across the three species. Black indicates low density variation and white high density variation. No point is shown for simulations in which an extinction occurred.

11.3 Variation in species densities

Another measure of ecological stability involves the change in species densities over time. The measurement is taken here by observing the species densities over a very long simulation, but rather than just recording the minimum density, the variation of each population from its mean density is recorded, and then these three variations are averaged out to give the values presented in figure 11.6. This is a measure of ecological constancy rather than persistence, because it quantifies the stability of the relative proportions of each species over time. It is equivalent to the 'average population variability' used

by Lawler (1993).

One motivation for using this measurement is that there are reasons for believing that density fluctuations can be expected to correlate positively with the maximum cluster sizes on the landscape; and these cluster sizes were found to correlate with time to extinction (see section 3.4.2, page 43). This positive correlation is expected because species tend to be more isolated in a community with bigger clusters, and when previously isolated groups come into contact with one another, it can lead to rapid changes in the state of quite large areas of the grid. Another way to think of this is to note that the sum of a small number of large changes, as found on a highly clustered grid, tends to vary more than the sum of a large number of small changes, which is what is found on a relatively unclustered grid. The combined effect on the densities over the entire grid of all these local changes is just the sum of what it going on locally, so the sum of the local variation should be visible at the larger scale in the form of bigger density fluctuations.

Figure 11.6 shows that when all three invasion rates are similar, the densities only change by a small amount, but when rates are uneven, densities fluctuate more wildly. For example, in the smallest variation observed in the simulations, the species densities varied by an average of less than 2%. This occurred with invasion rates at (0.35, 0.29, 0.36). But in the largest, most unstable variation, the three densities moved by an average of over 90%, with rates in the corner of the triangle at (0.1, 0.14, 0.76).

The importance of the ordering of invasion rates can be seen by the fact that darker points are not in phase with the outer triangle, but rotated slightly left as they are for TTE. The density variance over time, unlike the minimum density, differentiates simulations with few local extinctions from those with many local extinctions.

The average variation measure is not constant across the equilibrium invasion rates of the restraint system, and in fact figure 11.7 shows that it is possible to find two $r_s:r_p$ ratios, such as $r_p/(r_s + r_p) = 0.35$ and $r_p/(r_s + r_p) = 0.75$, that have very similar minimum variation and also very different equilibrium variation, so the average density variation not only lacks consistency at the equilibrium rates of the restraint system, it doesn't even appear to be proportional to the minimum variation. With the invasion rates marked by the rightmost red point in figure 11.6, the 500 × 500 grid is so unstable that data cannot be obtained for these rates due to extinctions.



Figure 11.7: Average variation in the three species densities over time at equilibrium invasion rates (solid line) and at the invasion rates that minimise average variation (dotted line) for $15 r_s:r_p$ ratios.

11.4 Cluster size

Because of the connection between clustering and stability, it is worth examining a measure of cluster size directly. Clustering is not strictly a kind of temporal stability (it's a static characteristic of the system), but nevertheless it turns out to be very well correlated with the measure of density variation over time.

The subcommunities envisaged in the community level selection interpretation of the restraint system are supposed to be 'emergent' subcommunities, created through spatial self-structuring (Johnson and Boerlijst, 2002). For this reason the small grids examined in section 11.1 will in most cases be unrepresentative of the J&S subcommunities because fixed 30×30 grids do not vary with respect to the quality of the self-structuring. Unfortunately, neither Johnson and Seinen (2002) nor Johnson and Boerlijst (2002) provides a suggestion for defining the extent of the emergent subcommunities, but it is reasonable to assume that the size of the subcommunities correlates with the spatial clustering on the grid.

It is easy to measure the level of clustering in a SCA by examining grids that have reached equilibrium and counting neighbouring sites that are in the same state. The details of the measurement used here are described in the caption to figure 11.8. It is the same measurement as the one used in section 3.4.2, but with a larger grid.

In the simulations summarised in figure 11.8, the average clustering measure



Figure 11.8: Comparison of the invasion rates that produce minimum clustering (black line) with the equilibrium invasion rates when r_r evolves freely (red line). The shading in each circle represents the average degree of clustering for one combination of invasion rates. For each rate combination, clustering was measured by running ten simulations for 12,000 generations each, on 500×500 grids. On each of these 10 grids, 1000 distinct sites were chosen at random from all over the grid, and for each random site, the number of contiguous sites to its right occupied by the same species as the original site was counted. Although the shading should therefore represent an average across 10,000 sites, in practice some of the simulations were ignored in the average. Points are only shown when at least 2000 sample sites were available.

ranges from 3, for rates of (0.35, 0.29, 0.36), up to 170, for the rates (0.7, 0.21, 0.09), and the mean value over all rates is 11. The theoretical maximum is 500, which can only occur if a single species occupies a horizontal strip the width of the entire grid.

The shapes of the minimum and equilibrium clustering, shown in figure 11.9, are similar to the density variations, including the particularly high value at the equilibrium when r_p is much greater than r_s . Clustering varies more across the equilibrium than it does at its minimum, so r_r does not evolve to limit



Figure 11.9: Minimum clustering (dotted line) and equilibrium clustering (solid line), measured using the technique described in figure 11.8 for $15 r_s:r_p$ ratios.

clustering to any threshold level.

But the lack of evidence for moderation of clustering by r_r in the experiments presented here is not enough to disprove the J&S hypothesis. A thorough approach would take a wide range of methods for defining subcommunity size as a function of system clustering, and then examine the stability of those potential subcommunities in order to see if there is some measure of subcommunity stability that appears to be moderated by the evolution of r_r . Unfortunately, due to the enormous number of ways to define subcommunities, this is too large a task to attempt here.

11.5 Summary

Johnson and Seinen (2002) showed that when r_r is allowed to evolve while r_s and r_p remain fixed, rock restrains its competitiveness to a value significantly less than its theoretical maximum. They suggest that this restraint is due to two opposing processes of individual level and community level selection, and that the community level selection pulls the entire community towards greater persistence stability.

The examination of the invasion rates at equilibrium provided in this chapter does not show that the multilevel selection picture is wrong, but it fails to confirm the 'selection for subcommunity stability' explanation using four simple quantifications of ecological stability.

11.5. SUMMARY

This is firstly because the invasion rate of rock does not appear to act in such a way as to moderate naturally unstable communities to any degree. In particularly unstable communities, highly competitive individuals do not appear to die out as a result of their subcommunities collapsing, instead they seem to thrive, because invasion rates are especially high in these unstable communities.

Secondly, none of the stability measurements remain constant at the equilibrium invasion rates reached by the J&S system. Although I have not been able to exactly reproduce the notion of persistence stability with any of the measures given here, the high variability of all measures across the equilibrium suggests that a model in which the supercompetitors die off when they bring their local subcommunity over some threshold of instability may be too simple.

Two-process explanations of restraint oversimplify

In chapter 10, I argued that to describe the restraint as the product of a process of selection among subcommunities could be misleading, and that an understanding of the details of the spatial RPS system was necessary for an understanding of the restraint phenomenon. It was suggested that the restraint phenomenon might be better understood as a product of the tradeoff between long-term and short-term fitness of individuals from the evolving species rather than by applying multilevel selection theory.

The experiments in this chapter also suggest that the long-term versus shortterm explanation of restraint is dubious. The time to extinction and minimum species density statistics, for example, measure the long-term fitness of the evolving species just as well as they measure system stability, because the evolving species is always one of the ones that goes extinct for all rates near the equilibrium. The fact that neither of these measurements settles at a consistent level is not only evidence for the lack of an instability threshold, it's also evidence for the lack of a 'long-term individual benefit' threshold. I suspect that attempts to reduce the essence of the restraint phenomenon to any simple model of two opposing selection processes will fail to capture some of the important details, regardless of whether one of the processes is described as higher-level selection or long-run individual benefit.

Chapter 12

Conclusions

The behaviour of the rock-paper-scissors system is highly dependent on the way in which the population is spatially structured. A SCA model in which individuals interact within a local neighbourhood always exhibits a high turnover of individuals in any particular region, but the overall densities of the three species become very stable. In a spatially-unstructured, infinite population model, densities oscillate forever with the same amplitudes, but in a similar finite-population model, the amplitudes increase over time until one species goes extinct.

In the first part of this thesis I have shown that in a model using continuous space, the resulting dynamics is very close to that of the lattice-based model, as long as interactions are localised. In both models, the species densities oscillate initially in similar patterns, stabilise to similar levels, and show similar spatial patterns on the landscape. The grid model accurately captures the essential features of the spatial system, and is computationally more efficient, so little is gained by using a continuous space model when interactions are known to be local.

A grid model with random, non-local interactions is just a pool of individuals without spatial structure, and in the case of RPS always ends in extinctions. But a continuous-space model does not have this property. As long as individuals have a location and a size, long-range dispersal is not the same thing as random interactions. It turns out that even a small amount of non-uniformity in the process by which individuals select others to interact with can significantly stabilise the species densities so that extinctions become very unlikely. The extent to which interactions deviate from complete uniformity is determined by the relative timescales of movement and interaction. Exactly how these two effects combine to bring about changes in density stability is not well understood, but could be discovered using a model with parameters for both speed of movement and frequency of interaction.

I have also shown the effect on system stability of gradual changes in the dispersal distance using a grid model. What is important is the distance over which interactions can take place relative to the total size of the world. The time until an extinction occurs increases as the size of the grid is increased, and also increases as the dispersal distance, or size of the interaction neighbourhood, is increased. The time to extinction in the RPS system, with its inherent stability, increases much faster than a system with only neutral evolution and no intransitivity. At a characteristic grid size, or neighbourhood size, the increase in stability becomes super-exponential. This means that for any particular maximum simulation length, there is effectively a threshold value for the grid size at which the stability jumps from a very low value to the maximum.

The neighbourhood size effect is slightly different to the grid size effect; when the grid is large there is an apparent threshold, as the neighbourhood size is decreased, at which stability suddenly jumps. But there are some grid sizes for which stability goes through an increase and then a subsequent decrease as neighbourhood size is reduced. The importance of this result is that contrary to the often expressed view that more spatial structure implies greater ecosystem diversity, there are occasions when more structure does the opposite, and increases the probability of ecosystem collapse. I have not discovered the exact mechanism causing the decrease in RPS stability for these very small neighbourhoods, but I believe it is the result of significant-sized pockets of single species becoming surrounded by their predators, a situation which is much less likely with a small increase in the number of neighbours.

In the second part of the thesis I have applied three spatial approximations to the RPS system and evaluated their accuracy with comparisons to grid simulations. Firstly, the pair approximation is used to derive equations for the densities of neighbouring pairs of grid sites. I show that these equations fail to express the essential stability of the grid-based RPS model for most starting points.

The local structure approximation based on 2×2 blocks of grid sites does not suffer from the same problem; it predicts that the RPS system is stable with all three species represented. It is not clear exactly why the pair approximation should fail while the 2×2 approximation succeeds. The 2×2 approximation includes more spatial information, but in some cases, it underestimates the density of homogeneous pairs of sites to a greater degree than the pair approximation. It may be that a successful local structure approximation requires the ability to represent neighbourhoods containing all the possible species in the system. The study of another approximation based on 3×1 blocks may provide further evidence for or against this hypothesis.

The need to estimate the risk of extinction in a spatial RPS system motivates a third simple approximation, which forces stability on the equations derived using the mean field assumption. It is trivial to modify the mean field equations so that the density trajectories are dragged in towards the fixed point by varying degrees depending on a new parameter. However, this approach is actually worse than the mean field when it comes to estimating the extinction risk in a grid simulation with sites initialised at random. This is because randomised initial conditions lead to an initial divergence in the densities while the initial spatial clusters are formed. It may be possible to devise an ODE model which tracks the degree of clustering in the system in addition to the densities. Rather than a local structure approaches to clustering, it might be better for such a model to parameterise the distribution of clusters of various sizes needed, because local structure models, like the mean field, are unable to account for the initial divergence of densities observed in the agent-based models.

The two models described in the third part of the thesis are interesting because they are able to produce RPS (or a similar four-species intransitive competition system) with the simple addition of an evolvable individual trait to a spatially-structured two-species system. In the first, individuals receive a certain quantity of resources based on their species, and are able to apportion those resources into interspecific or intraspecific competitiveness. When starting conditions are relatively even, the two species' densities stabilise at half and half even when their initial 'resource levels' are uneven. An analysis of the system using non-spatially-structured simulations, and of a similar system of differential equations, shows that stability of the densities at half and half is due to regulation by negative feedbacks rather than being a product of the spatial population structure. The feedback occurs because intraspecific competition is selected for in the species with the highest density, and interspecific competition is selected for in the species with the lower density. Both these selective forces act to equalise the densities, and the regulating effect becomes more efficient as more variation in inter- and intraspecific competitiveness is represented in the system.

It is spatial structure, however, that is responsible for the evolutionary branching event that splits at least one species into two distinct morphs, one maximising interspecific competitiveness and the other maximising intraspecific competitiveness. Spatial population structure can be shown to be the determining factor using a simple model of only three grid cells: the invasion fitness of maximally inter- or intraspecific competitors is greater than the invasion fitness of intermediate types, as long as there there is clustering of similar individuals in the three cells.

Even though species densities at equilibrium tend to even out regardless of the species' inherent strengths, the viability of morphs does in fact depend on those strengths. When the species are equal, both species split into extremely interspecific and extremely intraspecific morphs, competing in a cycle of four. When one species is more than about 20% stronger than the other, only the stronger species splits, the weaker one becomes extremely interspecific, and the weaker species competes with the two morphs of the stronger one in a RPS cycle. The ratio of species strengths at which the change occurs is quite robust, depending only to a minor extent on the size of mutations.

In the second model, interspecific toxin-production is the trait, rather than inter- and intraspecific competition. Rather than stabilising at half-and-half densities, the toxin-production model can be stable with uneven densities, and the stability of these densities can be predicted using a mean-field version of the model in which changes in densities and average toxin-production levels are given by differential equations. However, the real spatial simulations predict equal densities in a slightly different place to the mean field version, and the location of the equal-density point cannot be accounted for simply by adding a little bit of clustering of same-species individuals to the mean-field equations. This is possibly because clustering of individuals with similar trait values is not considered in this analysis.

Although the model is very simple, there are many different types of toxinproduction distribution into which the system settles. In the majority of the space of toxin and species strengths, one species maximises or minimises its toxin production exclusively, and the other species settles into a state where its average toxin-production is at an intermediate level, but sometimes this average combines two groups producing maximum and minimum amounts of toxin, and sometimes all individuals are clustered around the average level. I have not been able to predict exactly what causes the split into two groups, except to note that it is to do with the amount of room in the trait space between two clusters of phenotypes. If the trait's maximum or minimum value is forcing two groups to be closer than some minimum separation in the trait-space, then the groups join together when the weaker of the two cannot invade the other species fast enough to get away from the stronger group. Both of these models are potential explanations for the evolution of cyclic competitive behaviour. However, the first is problematic because species in existing intransitive ecosystems, as far as we know, do not possess any traits corresponding to the intraspecific competitiveness trait of the model. The toxin-production model is a better fit because toxin-based competition is prevalent in many of the natural ecosystems with intransitive competition. The 'splitting' behaviour of the traits in these models are one mechanism by which such ecosystems may increase in diversity.

In the fourth part of the thesis I argue that the occurrence of competitive restraint in the RPS system should not be explained by an appeal to community level selection theory. In Johnson and Seinen's model, multispecies subcommunities compete with one another within a larger community, and the subcommunities with the most persistence stability defeat those that are less stable. A supercompetitor will create instability in its subcommunity, so although the supercompetitor may do well within its subcommunity, its subcommunity will be less persistent (less fit) in the battle between subcommunities.

The first objection to this explanation is just that all subcommunities on a CA grid are arbitrary in their extent. Even when there is an element of spatial self-structuring of individuals, the way in which individuals interact in the CA model necessarily implies that there are as many subcommunities as there are individuals. Secondly, while a supercompetitor from an unstable RPS subcommunity is unfit in the long term, due to the survival of the weakest rule, members of other species within the same subcommunity may in fact be fitter. For these reasons I argue that the notion of community level selection adds little if any understanding to competitive restraint in RPS communities. It would better to explain the phenomenon with a model that describes the long-term fitness of supercompetitors.

If restraint were the product of a process of selection for stable subcommunities, then we might expect to see evidence for this by measuring the overall community stability at and away from the equilibrium at which competition becomes restrained. However, measurements of four kinds of stability at the equilibrium points, and across the entire space of possible invasion rates, reveals no pattern. In particular, stability at the equilibrium points shows a large amount of variation depending on the invasion rates of the two fixed species, meaning that there is no consistent threshold of instability over which subcommunities do not survive.

But although these measurements provide no evidence for the community

level selection view, they do not disprove it. This is because spatial patterns are different for different combinations of the fixed invasion rates, and subcommunities that have been defined using spatial self-structuring will have different sizes at these rate combinations. It could then be argued that when there are lots of subcommunities on the grid, the effect of instability in a few of them is diluted when measuring the stability of the entire community, but when there are a smaller number of subcommunities, instability in a few of them will count for more in the measurement of community stability. Such an argument would be unconvincing unless subcommunities can be defined using a notion of spatial self-structuring that is independent of stability.

Appendix

The adaptive dynamics method

I describe in general terms the method used in Kisdi (1999) and Doebeli and Dieckmann (2000), which provide typical examples of AD models. The first step is to describe the ecological dynamics of the mean phenotype of each species. If each species *s* has a mean trait value x_s , this is done by formulating an expression for the change in the population density of the mean phenotype of *s* over time, $d\rho(x_s)/dt$ for each species *s* where $\rho(x_s)$ stands for the density of the phenotype x_s . The resulting equations are of the form

$$\frac{d\rho(x_s)}{dt} = r_s \cdot \rho(x_s) \cdot w_s(x_s, P, X)$$
(A.1)

where $w_s(x_s, P, X)$ is some species-specific fitness function which depends on the population densities *P* of every species and their mean phenotypes *X*, and where r_s is the intrinsic growth rate of species *s* not dependent on the mean phenotypes. The articles by Kisdi and Doebeli and Dieckmann specify particular fitness functions where I have written $w_s(x_s, P, X)$.

The ecologically stable population densities $P^*(X)$ are the points at which the $d\rho(x_s)/dt = 0$ for the set of phenotypes X. At these stable densities, the *invasion fitness* of a rare mutant is examined. The invasion fitness $f_s(y_s, X)$ of a rare mutant y_s (of species *s*) is just its initial per capita growth rate, which is similar to equation (A.1):

$$f_s(y_s, X) = r_s \cdot w_s(y_s, P^*(X), X).$$
 (A.2)

But because this is y_s 's *per capita* growth rate, it is not dependent on y_s 's density, and therefore there is no $\rho(y_s)$ corresponding to the $\rho(x_s)$ in equation (A.1). And because y_s is rare, its effect on its own fitness can be ignored, so the species-dependent fitness function w_s still only depends on the densities

P and mean phenotypes *X* of the residents, even though a new phenotype has been introduced.

The derivative of a rare mutant's invasion fitness with respect to its trait value, evaluated at the mean phenotype,

$$\frac{\partial f_s(y_s, X)}{\partial y_s}|_{y_s = x_s} \tag{A.3}$$

describes the way in which evolution drives the mean phenotype x_s of species s. The points at which this expression is equal to zero are called *singular points*, and some of these singular points are also evolutionary attractors. If the mutant's fitness is at a maximum at at singular point, then this point is an ordinary stable attractor, but if the mutant's fitness is at a minimum at the singular point, then it is an evolutionary branching point. At a branching point, mutants of with trait values higher or lower than the mean phenotype will do better than the mean phenotype.

Invasion fitness at unstable points

In section 6.1.4, it is noted that AD is not a useful model in a two-species pre-RPS ecosystem, because such a system is always unstable up to the point at which one species goes extinct.

However, it is in fact possible to evaluate the invasion fitness of rare mutants (and the evolutionary direction of mean phenotypes) at ecologically unstable points, i.e. when $P \neq P^*(X)$, but of course there are infinitely many of these points, and the extra variables will make the expression (A.3) too difficult to analyse for most w_s functions.

Bibliography

- Abranches, J., Morais, P. B., Rosa, C. A., Mendonça-Hagler, L. C., and Hagler, A. N. (1997). The incidence of killer activity and extracellular proteases in tropical yeast communities. *Candadian Journal of Microbiology*, 43:328–336.
- Adams, J., Kinney, T., Thompson, S., Rubin, L., and Helling, R. B. (1979). Frequency-dependent selection for plasmid-containing cells of *Escherichia coli*. *Genetics*, 91:627–637.
- Adams, T. H., Wieser, J. K., and Yu, J.-H. (1998). Asexual sporulation in *Aspergillus nidulans*. *Microbiology and Molecular Biology Reviews*, 62:35–54.
- Ates, R. (1989). Aggressive behaviour in corals. *Freshwater and Marine Aquarium*, 12:104–112.
- Axelrod, R. (1984). *The Evolution of Co-operation*. Penguin Books.
- Bakus, G. J., Targett, N. M., and Schulte, B. (1986). Chemical ecology of marine organisms: an overview. *Journal of Chemical Ecology*, 12:951–987.
- Bell, G. (1997). Selection: The Mechanism of Evolution. Chapman & Hall.
- Berdy, J. (1974). Recent developments of antibiotic research and classification of antibiotics according to the chemical structure. *Advances in Applied Microbiology*, 18:309–406.
- Bjedov, I., Tenaillon, O., Gerard, B., Souza, V., Denamur, E., Radman, M., Taddei, F., and Matic, I. (2003). Stress-induced mutageneis in bacteria. *Science*, 300:1404–1409.
- Bolnick, D. I. (2004). Can intraspecific competition drive disruptive selection? An experimental test in natural populations of sticklebacks. *Evolution*, 58:608–618.
- Bürger, R. (2002a). Additive genetic variation under intraspecific competition and stabilizing selection: a two-locus study. *Theoretical Population Biology*, 61:197–213.

- Bürger, R. (2002b). On a genetic model of intraspecific competition and stabilizing selection. *The American Naturalist*, 160:661–682.
- Buss, L. W. (1980). Competitive intransitivity and size-frequency distributions of interacting populations. *Proceedings of the National Academy of Sciences*, 77:5355–5359.
- Buss, L. W. and Jackson, J. B. C. (1979). Competitive networks: nontransitive competitive relationships in cryptic coral reef environments. *The American Naturalist*, 113:223–234.
- Caraco, T., Duryea, M. C., Glavanakov, S., Maniatty, W., and Szymanski, B. K. (2001). Host spatial heterogeneity and the spread of vector-borne infection. *Theoretical Population Biology*, 59:185–206.
- Chesson, P. (2000). Mechanisms of maintenance of species diversity. *Annual Review of Ecology and Systematics*, 31:343–366.
- Cohn, D. (1992). *Separating Formal Bonds from Practical Performance in Learning Systems*. PhD thesis, University of Washington.
- Collins, J. P. (2003). What can we learn from community genetics? *Ecology*, 84:574–577.
- Connell, J. H. and Sousa, W. P. (1983). On the evidence needed to judge ecological stability or persistence. *The American Naturalist*, 121:789–824.
- Cranefield, G. (2001). Investigation of spatial ecologies. Project report, School of Mathematics, Statistics and Computer Science, Victoria University of Wellington, New Zealand. Available at http://www.mcs.vuw.ac.nz/comp/graduates/archives/ honours/2001/cranefield-report.ps.
- Czárán, T. L., Hoekstra, R. F., and Pagie, L. (2002). Chemical warfare between microbes promotes biodiversity. *Proceedings of the National Academy of Sciences*, 99:786–790.
- Dawkins, R. (1976). The Selfish Gene. Oxford University Press.
- Day, T. and Young, K. A. (2004). Competitive and facilitative evolutionary diversification. *BioScience*, 54:101–109.
- Dayton, P. K., Currie, V., Gerrodette, T., Keller, B. D., Rosenthal, R., and Tresca, D. V. (1984). Patch dynamics and stability of some california kelp communities. *Ecological Monographs*, 54:253–289.

- de Aguiar, M. A. M., Rauch, E. M., and Bar-Yam, Y. (2004). Invasion and extinction in the mean field approximation for a spatial host-pathogen model. *Journal of Statistical Physics*, 114:1417–1451.
- Dick, G. and Whigham, P. (2005). The behaviour of genetic drift in a spatiallystructured evolutionary algorithm. 2005 IEEE Congress on Evolutionary Computation, pages 1855–1860.
- Dieckmann, U. (1997). Can adaptive dynamics invade? *Trends in Ecology & Evolution*, 12:128–131.
- Dieckmann, U. and Doebeli, M. (1999). On the origin of species by sympatric speciation. *Nature*, 400:354–357.
- Dieckmann, U., Law, R., and Metz, J. A. J., editors (2000). *The Geometry of Ecological Interactions: Simplifying Spatial Complexity*. Cambridge University Press.
- Doebeli, M. and Dieckmann, U. (2000). Evolutionary branching and sympatric speciation caused by different types of ecological interactions. *The American Naturalist*, 156:S77–S101.
- Doebeli, M. and Dieckmann, U. (2004). Adaptive dynamics of speciation: spatial structure. In Dieckmann, U., Doebeli, M., Metz, J. A. J., and Tautz, D., editors, *Adaptive Speciation*, pages 140–167. Cambridge University Press, Cambridge.
- Donalson, D. D. and Nisbet, R. M. (1999). Population dynamics and spatial scale: effects of system size on population persistence. *Ecology*, 80:2492–2507.
- Drake, J. W., Charlesworth, B., Charlesworth, D., and Crow, J. F. (1998). Rates of spontaneous mutation. *Genetics*, 148:1667–1686.
- Durrett, R. and Levin, S. A. (1994a). The importance of being discrete (and spatial). *Theoretical Population Biology*, 46:363–394.
- Durrett, R. and Levin, S. A. (1994b). Stochastic spatial models: a user's guide to ecological applications. *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences*, 343:329–350.
- Durrett, R. and Levin, S. A. (1997). Allelopathy in spatially distributed populations. *Journal of Theoretical Biology*, 185:165–171.

- Durrett, R. and Levin, S. A. (1998). Spatial aspects of interspecific competition. *Theoretical Population Biology*, 53:30–43.
- Frank, S. A. (1994). Spatial polymorphism of bacteriocins and other allelopathic traits. *Evolutionary Ecology*, 8:369–386.
- Frank, S. A. (1995). Mutual policing and repression of competition in the evolution of cooperative groups. *Nature*, 377:520–522.
- Frean, M. and Abraham, E. R. (2001). Rock-scissors-paper and the survival of the weakest. *Proceedings of the Royal Society of London Series B: Biological Sciences*, 268:1323–1327.
- Ghiselin, M. T. (1987). Species concepts, individuality, and objectivity. *Biology and Philosophy*, 2:127–143.
- Gilpin, M. E. (1975a). *Group Selection in Predator-Prey Communities*. Princeton University Press, Princeton, NJ.
- Gilpin, M. E. (1975b). Limit cycles in competition communities. *The American Naturalist*, 109:51–60.
- Goodnight, C. J. (1990a). Experimental studies of community evolution. I. The response to selection at the community level. *Evolution*, 44:1614–1624.
- Goodnight, C. J. (1990b). Experimental studies of community evolution. II. The ecological basis of the response to community selection. *Evolution*, 44:1625–1636.
- Gordon, D. M. and Riley, M. A. (1999). A theoretical and empirical investigation of the invasion dynamics of colicinogeny. *Microbiology*, 145:655–661.
- Hamilton, W. D. (1964). The genetical evolution of social behaviour. *Journal* of *Theoretical Biology*, 7:1–52.
- Hamilton, W. D. (1975). Innate social aptitudes in man, an approach from evolutionary genetics. In Fox, R., editor, *Biosocial Anthropology*. Malaby.
- Harada, Y., Ezoe, H., Iwasa, Y., Matsuda, H., and Sato, K. (1995). Population persistence and spatially limited social interaction. *Theoretical Population Biology*, 48:65–91.
- Hassell, M. P., Comins, H. N., and May, R. M. (1994). Species coexistence and self-organizing spatial dynamics. *Nature*, 370:290–292.

- Hassell, M. P. and May, R. M. (1973). Stability in insect host-parasite models. *Journal of Animal Ecology*, 42:693–726.
- Hiebeler, D. (1997). Stochastic spatial models: From simulations to mean field and local structure approximations. *Journal of Theoretical Biology*, 187:307– 319.
- Hiebeler, D. (2000). Populations on fragmented landscapes with spatially structured heterogeneities: landscape generation and local dispersal. *Ecology*, 81(6):1629–1641.
- Hofbauer, J. and Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge University Press.
- Hull, D. L. (1988). Science as a Process. Chicago: University of Chicago Press.
- Ives, A. R., Turner, M. G., and Pearson, S. M. (1998). Local explanations of landscape patterns: can analytic approaches approximate simulation models of spatial processes? *Ecosystems*, 1:35–51.
- Iwasa, Y., Nakamaru, M., and Levin, S. A. (1998). Allelopathy of bacteria in a lattice population: competition between colicin-sensitive and colicinproducing strains. *Evolutionary Ecology*, 12:785–802.
- Johnson, C. R. (1997). Self-organising in spatial competition systems. In Klomp, N. I. and Lunt, I. D., editors, *Frontiers in ecology — building the links*, pages 245–263. Amsterdam, The Netherlands: Elsevier.
- Johnson, C. R. and Boerlijst, M. C. (2002). Selection at the level of the community: the importance of spatial structure. *Trends in Ecology & Evolution*, 17:83–90.
- Johnson, C. R. and Mann, K. H. (1988). Diversity, patterns of adaptation, and stability of nova scotian kelp beds. *Ecological Monographs*, 58:129–154.
- Johnson, C. R. and Seinen, I. (2002). Selection for restraint in competitive ability in spatial competition systems. *Proceedings of the Royal Society of London Series B: Biological Sciences*, 269:655–663.
- Kerr, B., Riley, M. A., Feldman, M. W., and Bohannan, B. J. M. (2002). Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors. *Nature*, 418:171–174.
- Kirkup, B. C. and Riley, M. A. (2004). Antibiotic-mediated antagonism leads to a bacterial game of rock-paper-scissors *in vivo*. *Nature*, 428:412–414.

- Kisdi, E. (1999). Evolutionary branching under asymmetric competition. *Journal of Theoretical Biology*, 197:149–162.
- Kisdi, E. and Geritz, S. A. H. (1999). Adaptive dynamics in allele space: Evolution of genetic polymorphism by small mutations in a heterogeneous environment. *Evolution*, 53:993–1008.
- Kumar, S. and Subramanian, S. (2002). Mutation rates in mammalian genomes. *Proceedings of the National Academy of Sciences*, 99:803–808.
- Lavin, P. A. and McPhail, J. D. (1985). The evolution of freshwater diversity in the threespine stickleback (*Gasterosteus aculeatus*): site-specific differentiation of trophic morphology. *Canadian Journal of Zoology*, 83:2632–2638.
- Lawler, S. P. (1993). Species richness, species composition and population dynamics of protists in experimental microcosms. *Journal of Animal Ecology*, 62:711–719.
- Leigh, E. G. J. (1994). Do insect pests promote mutualism among tropical trees? *Journal of Ecology*, 82:677–680.
- Mansfield, R. (2001). The evolution of policing. Project report, School of Mathematics, Statistics and Computer Science, Victoria University of Wellington, New Zealand. Available at http://www.mcs.vuw.ac.nz/comp/graduates/archives/honours/2000/Mansfield-report.ps.
- Matsuda, H., Ogita, N., Sasati, A., and Sato, K. (1992). Statistical mechanics of population: the lattice Lotka-Volterra model. *Progress of Theoretical Physics*, 88:1035–1049.
- Matsuda, H., Tamachi, N., Sasaki, A., and Ogita, N. (1987). A lattice model for population biology. In Teramoto, E. and Yamaguti, M., editors, *Mathematical Topics in Biology, Morphogenesis and Neurosciences*, pages 154–161. Springer Lecture Notes in Biomathematics 71.
- May, R. M. and Leonard, W. (1975). Nonlinear aspects of competition between three species. *SIAM Journal on Applied Mathematics*, 29:243–253.
- Maynard Smith, J. (1964). Group selection and kin selection. *Nature*, 201:1145–1147.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge University Press.

- Maynard Smith, J. and Szathmáry, E. (1995). *The Major Transitions in Evolution*. W. H. Freeman.
- Metz, J. A. J., Geritz, S. A. H., Meszéna, G., Jacobs, F. J. A., and van Heerwaarden, J. S. (1996). Adaptive dynamics, a geometrical study of the consequences of nearly faithful reproduction. In van Strien, S. J. and Verduyn Lunel, S. M., editors, *Stochastic and Spatial Structures of Dynamical Systems*, pages 183–231. North Holland, Amsterdam.
- Miller, J. F. (1996). Spontaneous mutators in bacteria: insights into pathways of mutagenesis and repair. *Annual Review of Microbiology*, 50:625–643.
- Mitteldorf, J., Croll, D. H., and Ravela, S. C. (2002). Multilevel selection and the evolution of predatory restraint. In Standish, R. K., Abbass, H. A., and Bedau, M. A., editors, *Artificial Life VIII*, pages 146–152. MIT Press.
- Mitteldorf, J. and Wilson, D. S. (2000). Population viscosity and the evolution of altruism. *Journal of Theoretical Biology*, 204:481–496.
- Moxon, E. R., Rainey, P. B., Nowak, M. A., and Lenski, R. E. (1994). Adaptive evolution of highly mutable loci in pathogenic bacteria. *Current Biology*, 4:24–33.
- Neuhauser, C., Andow, D. A., Heimpel, G. E., May, G., Shaw, R. G., and Wagenius, S. (2003). Community genetics: expanding the synthesis of ecology and genetics. *Ecology*, 84:545–558.
- Nowak, M. A. and May, R. M. (1992). Evolutionary games and spatial chaos. *Nature*, 359:826–829.
- Pagie, L. and Hogeweg, P. (1999). Colicin diversity: a result of ecoevolutionary dynamics. *Journal of Theoretical Biology*, 196:251–261.
- Poole, A. M., Phillips, M. J., and Penny, D. (2003). Prokaryote and eukaryote evolvability. *BioSystems*, 69:163–185.
- Price, G. R. (1970). Selection and Covariance. Nature, 277:520–521.
- Queller, D. C. (1992). Does population viscosity promote kin selection? *Trends in Ecology & Evolution*, 7:322–324.
- Rand, D. A. (1999). Correlation equations and pair approximations for spatial ecologies. *CWI Quarterly*, 12:329–368.
- Reeves, P. (1972). The Bacteriocins. Springer Verlag, New York.

- Ridgeway, M. S. and McPhail, J. D. (1984). Ecology and evolution of sympatric sticklebacks (*Gasterosteus*): Mate choice and reproductive isolation in the enos lake species pair. *Canadian Journal of Zoology*, 62:1813–1818.
- Riley, M. A. (1998). Molecular mechanisms of bacteriocin evolution. *Annual Review of Genetics*, 32:255–278.
- Riley, M. A. and Gordon, D. M. (1996). The ecology and evolution of bacteriocins. *Journal of Industrial Microbiology*, 17:151–158.
- Sato, K. and Iwasa, Y. (2000). Pair approximations for lattice-based ecological models. In Dieckmann, U., Law, R., and Metz, J. A. J., editors, *The Geometry* of Spatial Interactions: Simplifying Spatial Complexity, pages 341–358. Cambridge University Press.
- Schluter, D. and McPhail, J. D. (1992). Ecological character displacement and speciation in sticklebacks. *The American Naturalist*, 140:85–108.
- Seger, J. (1985). Intraspecific resource competition as a cause of sympatric speciation. In Greenwood, P. J., Harvey, P. H., and Slatkin, M., editors, *Evolution: Essays in honour of John Maynard Smith*, pages 43–53. Cambridge University Press, Cambridge.
- Silvertown, J., Holtier, S., Johnson, S., and Dale, P. (1992). Cellular automaton models of interspecific competition for space the effect of pattern on process. *Journal of Ecology*, 80:527–534.
- Sinervo, B. and Lively, C. M. (1996). The rock-paper-scissors game and the evolution of alternative male strategies. *Nature*, 340:240–243.
- Sober, E. and Wilson, D. S. (1998). *Unto Others: the Evolution and Psychology of Unselfish Behaviour*. Harvard University Press.
- Starmer, W. T., Ganter, P. F., Aberdeen, V., Lachance, M.-A., and Phaff, H. J. (1987). The ecological role of killer yeasts in natural communities of yeasts. *Candadian Journal of Microbiology*, 33:783–796.
- Stone, L. and Roberts, A. (1991). Conditions for a species to gain advantage from the presence of competitors. *Ecology*, 72:1964–1972.
- Suzuki, D. T., Griffiths, A. J. F., Miller, J. H., and Lewontin, R. C. (1989). *An Introduction to Genetic Analysis*. W. H. Freeman.
- Szabó, G. and Hauert, C. (2002). Evolutionary prisoner's dilemma games with voluntary participation. *Physical Review E*, 66:062903.

- Szabó, G., Szolnoki, A., and Izsák, R. (2004). Rock-scissors-paper game on regular small-world networks. *Journal of Physics A: Mathematical and General*, 37:2599–2609.
- Taylor, P. D. (1992). Altruism in viscous populations an inclusive fitness model. *Evolutionary Ecology*, 6:352–356.
- Tilman, D. and Kareiva, P., editors (1997). *Spatial Ecology: The Role of Space in Population Dynamics and Interspecific Interactions*. Princeton University Press.
- van Baalen, M. (2000). Pair approximations for different spatial geometries. In Dieckmann, U., Law, R., and Metz, J. A. J., editors, *The Geometry of Spatial Interactions: Simplifying Spatial Complexity*, pages 359–387. Cambridge University Press.
- van Baalen, M. and Rand, D. A. (1998). The unit of selection in viscous populations and the evolution of altruism. *Journal of Theoretical Biology*, 193:631–648.
- Whitham, T. G., Young, W. P., Martinsen, G. D., Gehring, C. A., Schweitzer, J. A., Shuster, S. M., Wimp, G. M., Fischer, D. G., Bailey, J. K., Lindroth, R. L., Woolbright, S., and Kuske, C. R. (2003). Community and ecosystem genetics: a consequence of the extended phenotype. *Ecology*, 84:559–573.
- Wickner, R. B. (1992). Double-stranded and single-stranded RNA viruses of *Saccharomyces cerevisiae*. *Annual Review of Microbiology*, 46:347–375.
- Williams, G. C. (1966). Adaptaion and Natural Selection: A Critique of Some Current Evolutionary Thought. Princeton: Princeton University Press.
- Williams, G. C. and Williams, D. C. (1957). Natural selection of individually harmful social adaptations among sibs with special reference to social insects. *Evolution*, 11:32–39.
- Wilson, D. S. (1997). Biological communities as functionally organized units. *Ecology*, 78:2018–2024.
- Wilson, D. S., Pollock, G. B., and Dugatkin, L. A. (1992). Can altruism evolve in purely viscous populations? *Evolutionary Ecology*, 6:331–341.
- Wilson, D. S. and Sober, E. (1994). Reintroducing group selection to the human behavioural sciences. *Behavioral and Brain Sciences*, 17:585–654.
- Wilson, D. S. and Swenson, W. (2003). Community genetics and community selection. *Ecology*, 84:586–588.

- Wolfe, K. H., Sharp, P. M., and Li, W.-H. (1989). Mutation rates differ among regions of the mammalian genome. *Nature*, 337:283–285.
- Wynne-Edwards, V. C. (1962). *Animal Dispersion in Relation to Social Behaviour*. Oliver & Boyd, Edinburgh.