

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-paper-scissors 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 rock-paper-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-paper-scissors 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	Rock-paper-scissors ecosystems	1
1.1	The mean field model	3
1.2	Stochastic cellular automaton model	4
1.3	Finite-population model with long-range dispersal	6
1.4	Outline of the thesis	7
I	The effects of spatial structure in agent-based models	13
2	RPS in continuous space	15
2.1	A continuous-space model for RPS	16
2.2	Comparison 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	Discussion	27
3	Intermediate-range dispersal	29
3.1	Time to extinction (TTE) experiments	30
3.2	Increasing 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	Decreasing 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	Combinations 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 neighbourhoods	43
3.5	Summary and discussion	44

II Approximations of spatial structure 47

4 Pair approximations for RPS ecosystems 49

4.1	Pair and local structure approximations	50
4.2	Application to RPS	51
4.2.1	Pair transformation dynamics	51
4.2.2	Symmetry in pair types	51
4.2.3	Counting of pairs	52
4.2.4	Homogeneous and heterogeneous pairs	53
4.3	Derivation of pair correlation equations	54
4.3.1	Rate of change of homogeneous pairs	54
4.3.2	Rate of change of heterogeneous pairs	55
4.3.3	Rates of transformation of pairs	55

4.3.4	The pair closure	56
4.3.5	Rate of change of singleton densities	58
4.4	Failure 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 2×2 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 ' 2×2 block' approximation	70
4.6	Behaviour of the 2×2 model	71
4.6.1	Stability	72
4.6.2	Ecosystem composition	72
4.6.3	Neighbour correlations and species clustering	74
4.7	Discussion	76
5	Direct spatial stability approximation	79
5.1	Explicit spatial stability	80
5.2	Behaviour of the approximation	80
5.2.1	Extinction risk	81
5.2.2	Predictions of extinction risk	82

III	The evolution of competitive cycles	85
6	Intraspecific competition	87
6.1	Modelling 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 intraspecific 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	Behaviour 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	Realistic 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 intraspecific competitiveness trait	113
6.6	Summary	115
7	Density stability in the intra model	117
7.1	A nonspatial 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	Properties of the nonspatial model	122
7.3	Nonspatial 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	Importance of spatial structure for coexistence	129
7.5	Summary	131
8	Disruptive selection in the intra model	133
8.1	Spatial interactions create local clustering	134
8.2	The three-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 fastest invaders survive	140
8.4	Asymmetric competition	142
8.5	Summary	142
9	Cyclic competition among toxin-producers	145

9.1	Modelling toxin-based competition	146
9.1.1	Theoretical models of microbial toxin systems	146
9.1.2	Advantages of the intraspecific model	149
9.1.3	Application of the intraspecific model	150
9.2	A two-species toxin-growth tradeoff model	151
9.3	Species densities	153
9.3.1	Coexistence and stability	154
9.3.2	Stability in spatially unstructured populations	154
9.3.3	Net invasion rates for pure strategies	156
9.3.4	Mean-field model	157
9.3.5	Local clustering of species	160
9.4	Distribution of toxin-producers	164
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	168
9.4.4	Summary of equilibrium behaviour	170
9.5	Invasion of an established population	171
9.6	Generalised toxin-growth model	172
9.6.1	Mean field equations	173
9.6.2	Species densities in the general model	173
9.6.3	Distribution of toxin-producers in the general model	176
9.6.4	Invasion of an established population	179
9.7	Summary	180

IV Competitive restraint 183

10 Restraint and community level selection 185

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	Population structure	189
10.2.1	What is a group?	189
10.2.2	What is a (sub)community?	189
10.3	Spatially 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 communities	192
10.4	Communities 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	Conclusion	197
11	Restraint and stability	199
11.1	Time to 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	Minimum species density	205
11.3	Variation in species densities	208

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		
		r	s	p
Player 1	r	0	1	-1
	s	-1	0	1
	p	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-paper-scissors, 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).

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:

$$\frac{d\rho_r}{dt} = \rho_r\rho_s r_r - \rho_p\rho_r r_p \quad (1.1)$$

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

$$\rho_r = \frac{r_s}{r_r + r_s + r_p}, \quad \rho_s = \frac{r_p}{r_r + r_s + r_p}, \quad \rho_p = \frac{r_r}{r_r + r_s + r_p}.$$

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 left-most 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, non-lattice-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 non-mobile 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

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 multi-species subcommunities, and that in the competition among subcommunities, those that are the most stable tend to win. They claim that evolving-species in-

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.