

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 multi-species 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.

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-paper-scissors 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 2*a* 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 between-community 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 sub-

communities 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-

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 subcommunity could end up being quite well represented in the 'children' of these 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 group-level 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

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.

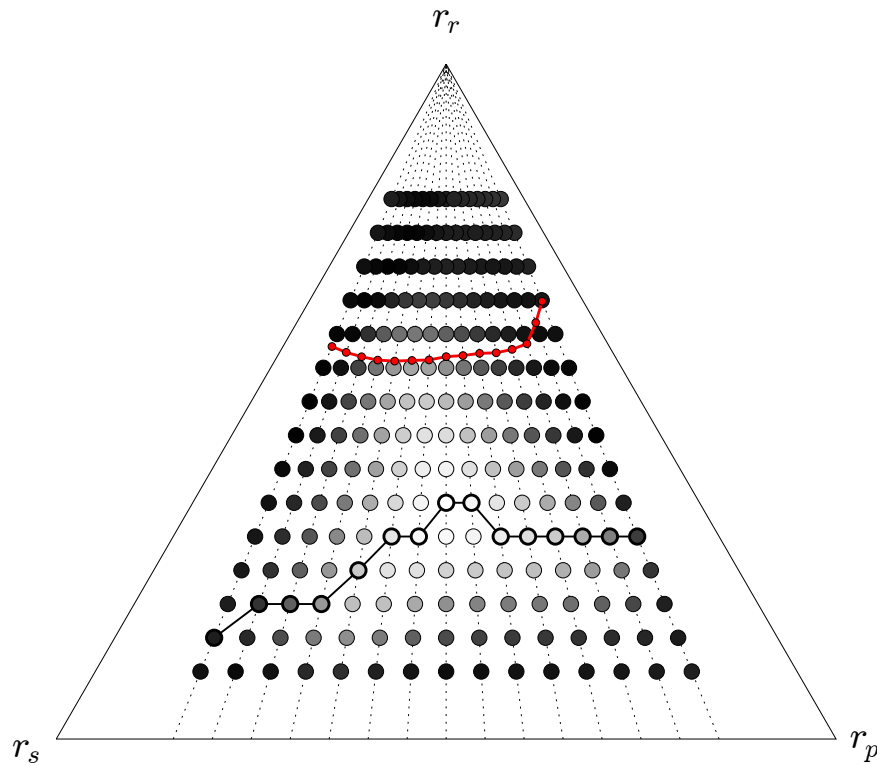


Figure 11.1: A comparison of the longest average time to extinction with the equilibrium points of the restraint system described in Johnson and Seinen (2002). Each dotted line shows one ratio of r_s to r_p , and the shaded circles on the dotted lines show the TTE (averaged over 2500 simulations on a 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 $r_r > r_p > r_s$).

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

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.

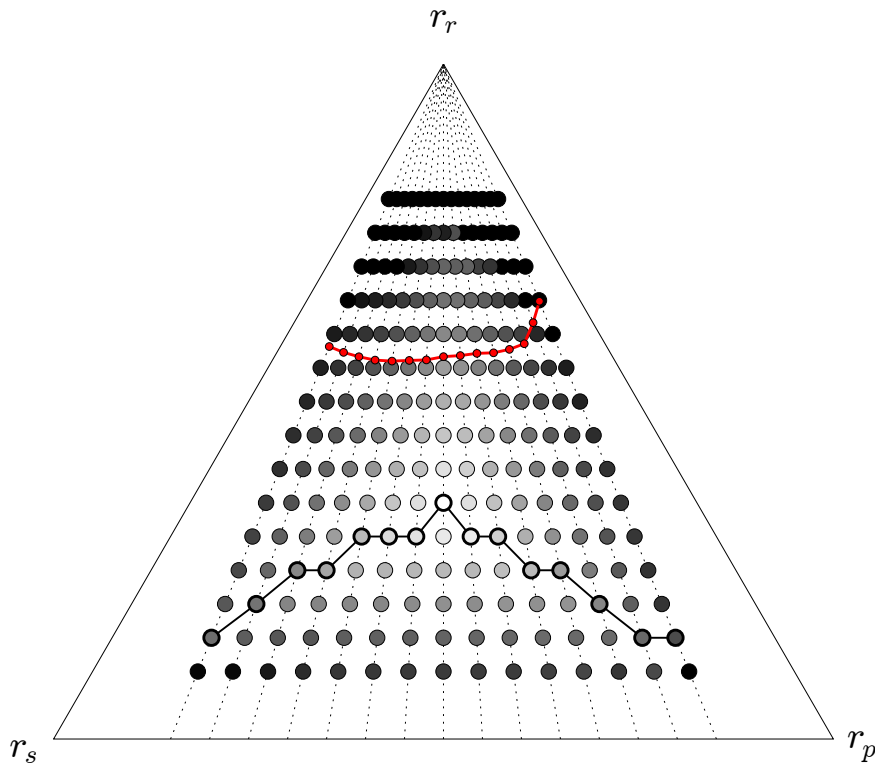


Figure 11.4: Minimum extinction risk (black line) compared against extinction risk at restraint equilibrium (red line). The shading in a circle indicates the minimum density of the least common species during the final 12,000 generations of a 24,000 generation RPS simulation on a 500×500 grid for a fixed 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 predated 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.

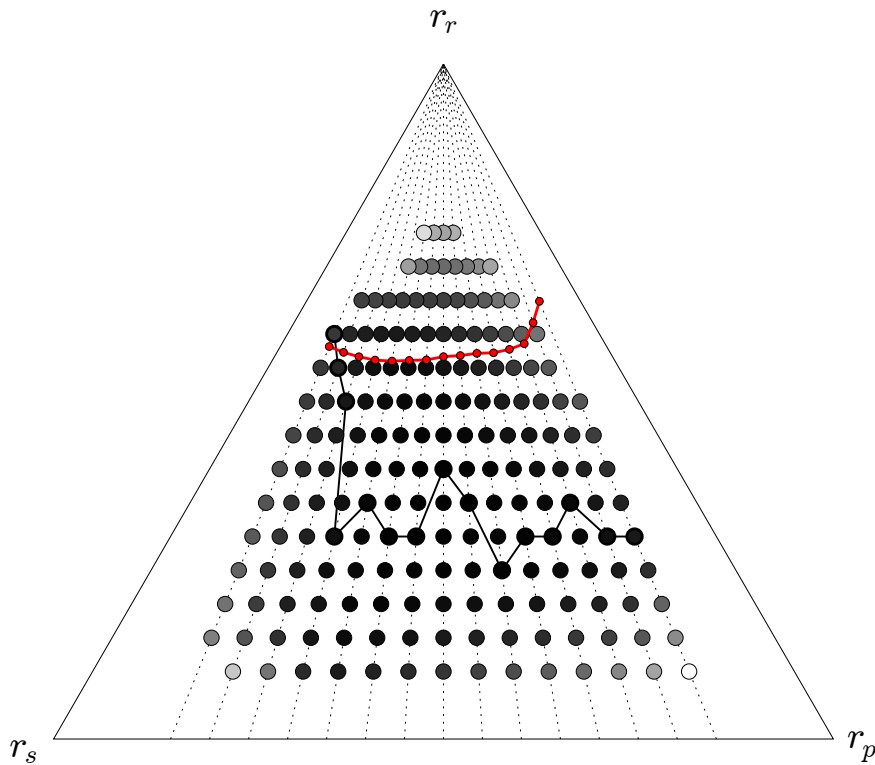


Figure 11.6: Comparison of the invasion rates which produce minimum density variation over time (black line) with the equilibrium invasion rates when r_r evolves freely (red line). The shading in the circles represents the average, over all three species, of the amount by which their densities fluctuate 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 is 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.

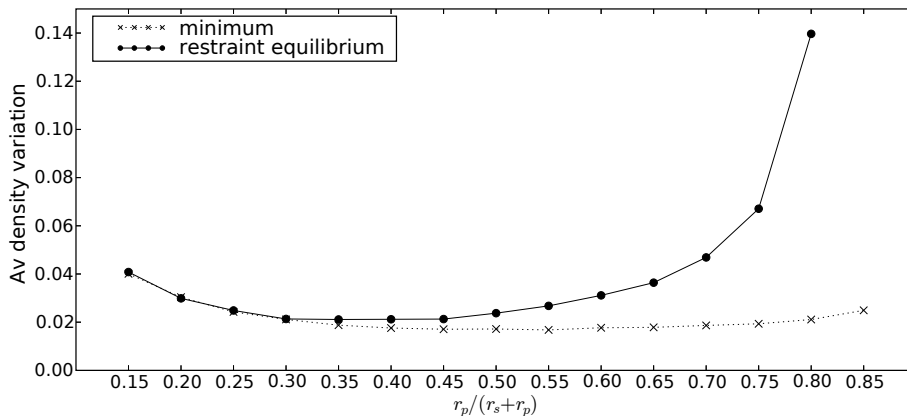


Figure 11.7: Average variation in the three species densities over time at equilibrium invasion rates (solid line) and at the invasion rates that minimise average variation (dotted line) for 15 $r_s:r_p$ ratios.

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 reached extinction during the 12,000 timesteps, and sites from these 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

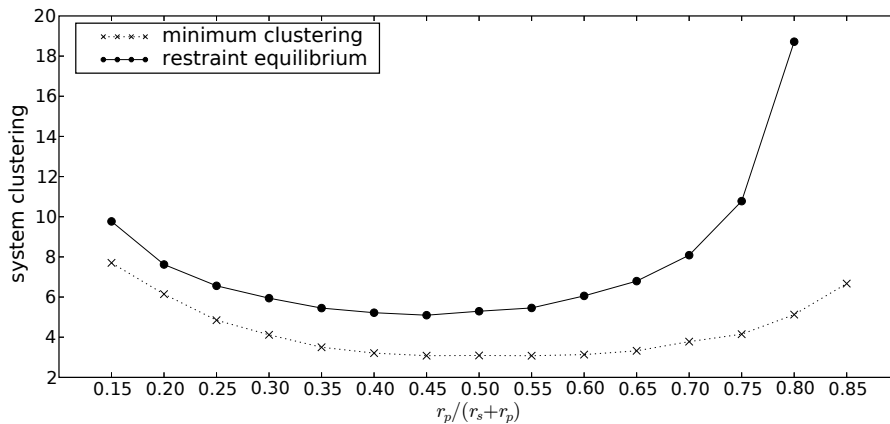


Figure 11.9: Minimum clustering (dotted line) and equilibrium clustering (solid line), measured using the technique described in figure 11.8 for 15 $r_s:r_p$ ratios.

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.

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 short-term 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 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 toxin-production 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

$$\frac{d\rho(x_s)}{dt} = r_s \cdot \rho(x_s) \cdot w_s(x_s, P, X) \quad (\text{A.1})$$

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):

$$f_s(y_s, X) = r_s \cdot w_s(y_s, P^*(X), X). \quad (\text{A.2})$$

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,

$$\frac{\partial f_s(y_s, X)}{\partial y_s} \Big|_{y_s=x_s} \quad (\text{A.3})$$

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 a 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. *Canadian 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., Tenailon, O., Gerard, B., Souza, V., Denamur, E., Radman, M., Taddei, F., and Matic, I. (2003). Stress-induced mutagenesis 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 californian 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 spatially-structured 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 colicin-producing 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 Wageningen, 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 eco-evolutionary 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*, 227:520–521.
- 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. *Canadian 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). *Adaptation 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.