

Part II

Approximations of spatial structure

Chapter 4

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

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

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

4.1 Pair and local structure approximations

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

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

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

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

4.2 Application to RPS

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

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

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

4.2.1 Pair transformation dynamics

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

4.2.2 Symmetry in pair types

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

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





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

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

4.2.3 Counting of pairs

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

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



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

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

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

4.2.4 Homogeneous and heterogeneous pairs

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

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

4.3 Derivation of pair correlation equations

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

4.3.1 Rate of change of homogeneous pairs

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

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

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

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

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

$$\frac{d\rho_{aa}}{dt} = R[ab \rightarrow aa] - R[aa \rightarrow ca] \quad (4.1)$$

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

$$\begin{aligned} \text{prey}(r) &= s, & \text{prey}(s) &= p, & \text{prey}(p) &= r, \text{ and} \\ \text{predator}(r) &= p, & \text{predator}(s) &= r, & \text{predator}(p) &= s. \end{aligned}$$

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

4.3.2 Rate of change of heterogeneous pairs

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

$$\frac{d\rho_{ab}}{dt} = R[bb \rightarrow ab] + R[ca \rightarrow ab] - R[ab \rightarrow bc] - R[ab \rightarrow aa] \quad (4.2)$$

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

4.3.3 Rates of transformation of pairs

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

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

Case 1: $ab \rightarrow aa$

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

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

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

$$R[rs \rightarrow rr] = \rho_{rs}r_r(1 + (Q - 1)q_{r|sr}). \quad (4.3)$$

Case 2: $aa \rightarrow ca$

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

$$R[rr \rightarrow pr] = \rho_{pr}r_p(Q - 1)q_{r|rp}. \quad (4.4)$$


Case 3: $ca \rightarrow ab$

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

$$R[pr \rightarrow rs] = \rho_{sp}r_s(Q - 1)q_{r|ps}. \quad (4.5)$$

4.3.4 The pair closure

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



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

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

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

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

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

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

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

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

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

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

$$\begin{aligned} R[ab \rightarrow aa] &= \rho_{ab} r_a (1 + (Q - 1) q_{a|b}) \\ R[aa \rightarrow ca] &= \rho_{ca} r_c (Q - 1) q_{a|a} \\ R[ca \rightarrow ab] &= \rho_{bc} r_b (Q - 1) q_{a|c} \end{aligned} \quad (4.7)$$

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

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

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

4.3.5 Rate of change of singleton densities

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

$$\frac{d\rho_r}{dt} = \frac{d\rho_{rr}}{dt} + \frac{1}{2} \cdot \frac{d\rho_{rs}}{dt} + \frac{1}{2} \cdot \frac{d\rho_{pr}}{dt}$$

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

$$\begin{aligned} \frac{d\rho_r}{dt} &= R[rs \rightarrow rr] - R[rr \rightarrow pr] \\ &\quad + \frac{1}{2} (R[ss \rightarrow rs] + R[pr \rightarrow rs] - R[rs \rightarrow rr] - R[rs \rightarrow sp]) \\ &\quad + \frac{1}{2} (R[rr \rightarrow pr] + R[sp \rightarrow pr] - R[pr \rightarrow pp] - R[pr \rightarrow rs]) \\ &= \frac{1}{2} (R[rs \rightarrow rr] + R[ss \rightarrow rs] + R[sp \rightarrow pr] \\ &\quad - R[rr \rightarrow pr] - R[rs \rightarrow sp] - R[pr \rightarrow pp]) \\ &= \frac{1}{2} (\rho_{rs}r_r(1 + (Q-1)q_{r|s}) + \rho_{rs}r_r(Q-1)q_{s|s} + \rho_{rs}r_r(Q-1)q_{p|s} \\ &\quad - \rho_{pr}r_p(Q-1)q_{r|r} - \rho_{pr}r_p(Q-1)q_{s|r} - \rho_{pr}r_p(1 + (Q-1)q_{p|r})) \\ &= \frac{1}{2} (\rho_{rs}r_r(1 + (Q-1)(q_{r|s} + q_{s|s} + q_{p|s})) \\ &\quad - \rho_{pr}r_p(1 + (Q-1)(q_{r|r} + q_{s|r} + q_{p|r}))). \end{aligned}$$

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

$$\frac{d\rho_r}{dt} = \frac{Q}{2} (\rho_{rs}r_r - \rho_{pr}r_p). \quad (4.8)$$

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

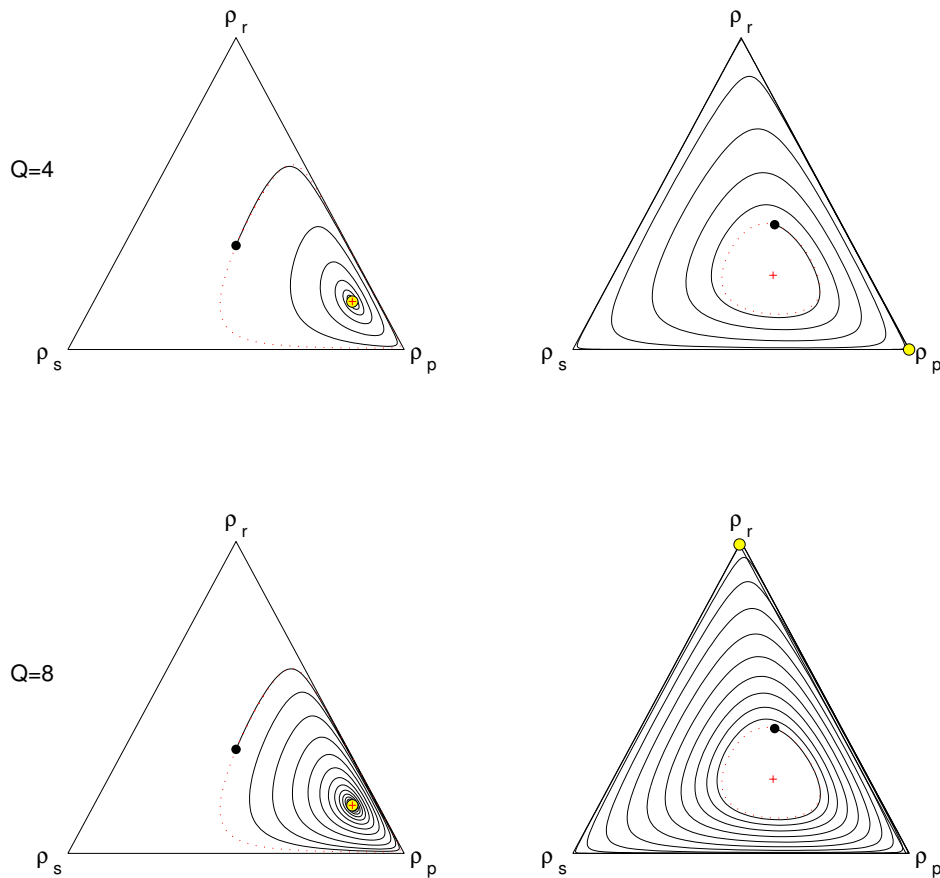


Figure 4.2: Examples of trajectories of species densities for the RPS pair approximation equations given in section 4.3. The black dot shows the initial singleton densities, the black line the trajectory of the densities over time, and the yellow dot the final stable densities. The red cross and dotted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

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

4.4 Failure of the pair approximation for RPS

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

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

4.4.1 Importance of neighbourhood size

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

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

4.4.2 Extinctions more common than coexistence

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

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

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

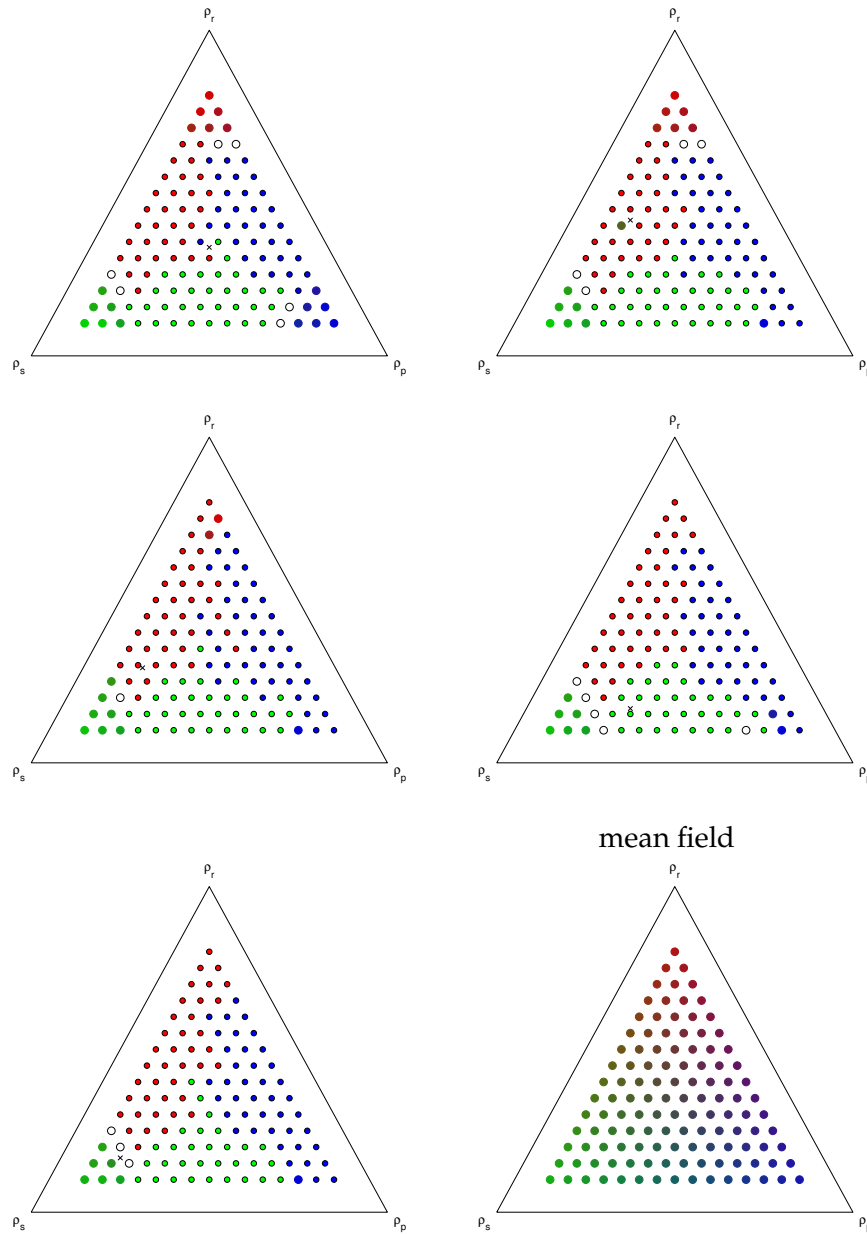


Figure 4.3: Equilibrium species densities for the RPS pair approximation system with $Q = 4$. Each circle represents a final species density. Small circles indicate extinctions with the winning species indicated by the colour of the circle (where rock = red, scissors = green, paper = blue). The larger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

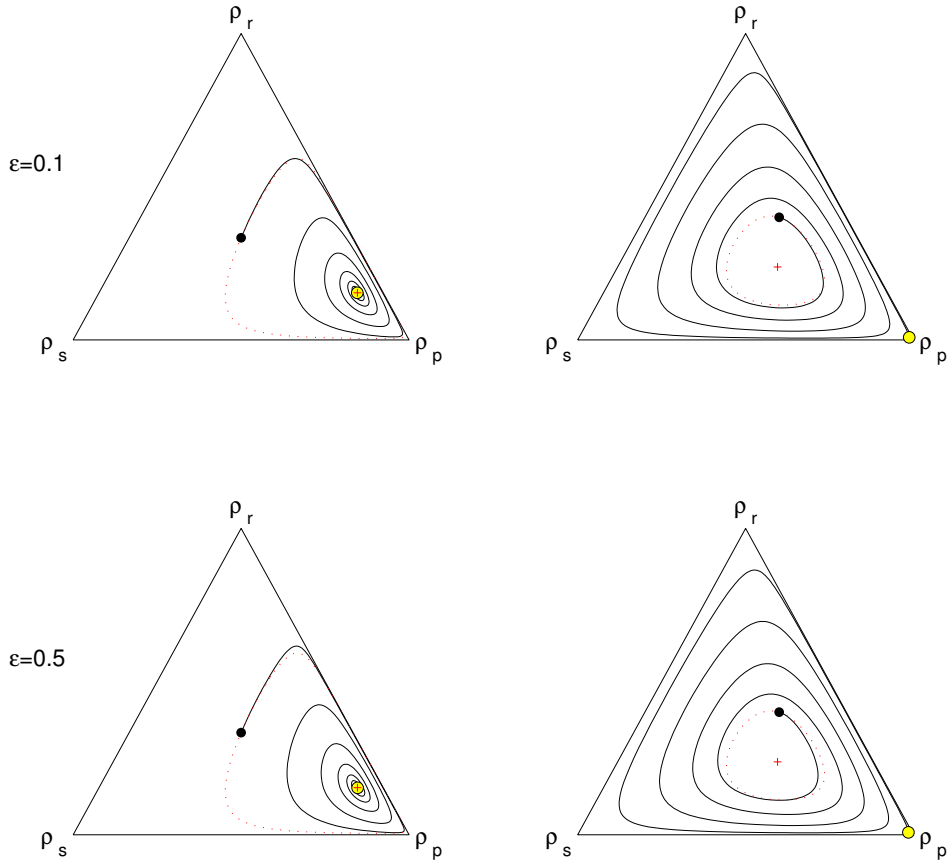


Figure 4.4: Examples of the pair approximation trajectories with $Q = 4$ when a correlation quantified by ε is added to the initial pair densities. The resulting trajectories initially deviate slightly from those with uncorrelated initial pair densities in figure 4.2. The final densities are unaffected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

4.4.3 Effects of initial pair densities

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

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

$$\begin{aligned} \rho_{rs} + \rho_{sp} + \rho_{pr} &= (1 - \varepsilon)(2\rho_r \rho_s + 2\rho_s \rho_p + 2\rho_p \rho_r), \\ \rho_{rr} + \rho_{ss} + \rho_{pp} &= \rho_r^2 + \rho_s^2 + \rho_p^2 + \varepsilon(2\rho_r \rho_s + 2\rho_s \rho_p + 2\rho_p \rho_r). \end{aligned}$$

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

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

$$\rho'_{rr} = \rho_r^2 \left(1 + \frac{\varepsilon(2\rho_r\rho_s + 2\rho_s\rho_p + 2\rho_p\rho_r)}{\rho_r^2 + \rho_s^2 + \rho_p^2} \right),$$

$$\rho'_{rs} = 2\rho_r\rho_s(1 - \varepsilon).$$

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

4.4.4 Pair approximation compared to SCA simulation

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

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

4.4.5 Summary

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

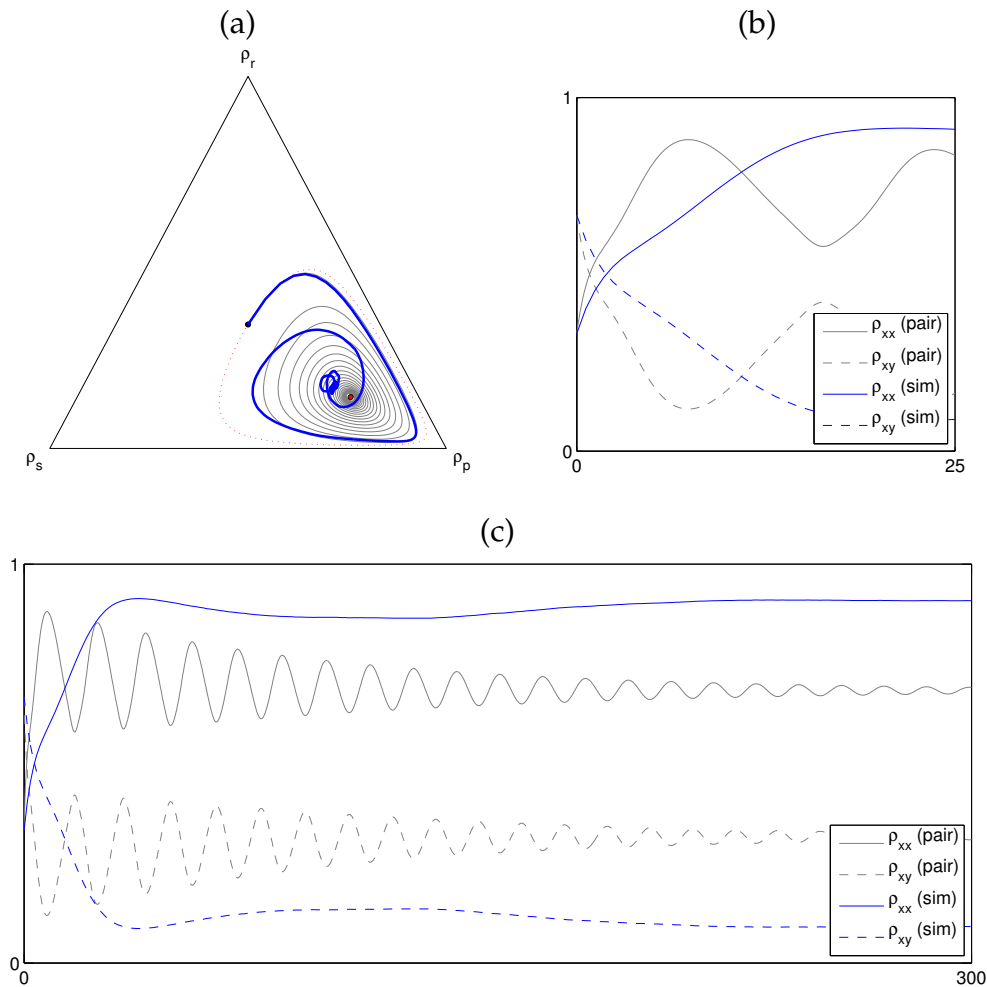


Figure 4.5: Comparison of the pair approximation to the SCA simulation. (a) Trajectories of the mean field (dotted red line), pair approximation (grey line) and 700×700 SCA simulation (blue line), for equal initial densities and invasion rates (1,0.2,0.25). (b) and (c) show the pair approximation's predicted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

4.5 2×2 approximations

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




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

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

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

4.5.1 RPS and 2×2 blocks

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

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

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

$$\rho_{\begin{bmatrix} xx \\ yz \end{bmatrix}} = \rho_{\begin{bmatrix} xx \\ zy \end{bmatrix}}, \quad x, y, z \in \{r, s, p\}. \quad (4.10)$$

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

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

4.5.2 Block densities, ordered and unordered

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

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

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

$$\rho_{pr} = \rho_{\begin{bmatrix} pr \\ rr \end{bmatrix}} + \rho_{\begin{bmatrix} rp \\ rr \end{bmatrix}} + \rho_{\begin{bmatrix} rr \\ rp \end{bmatrix}} + \rho_{\begin{bmatrix} rr \\ pr \end{bmatrix}},$$

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

$$\rho_{pr} = \rho_{\begin{bmatrix} pr \\ sr \end{bmatrix}} + \rho_{\begin{bmatrix} sp \\ rr \end{bmatrix}} + \rho_{\begin{bmatrix} rs \\ rr \end{bmatrix}} + \rho_{\begin{bmatrix} rr \\ ps \end{bmatrix}} + \rho_{\begin{bmatrix} sr \\ pr \end{bmatrix}} + \rho_{\begin{bmatrix} ps \\ rr \end{bmatrix}} + \rho_{\begin{bmatrix} rp \\ rs \end{bmatrix}} + \rho_{\begin{bmatrix} rr \\ sp \end{bmatrix}}.$$

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

$$\begin{aligned}
\text{ordered}_{aa}^{aa} &= \{[aa]\} \\
\text{ordered}_{aa}^{ca} &= \{[ca], [aa], [ac], [ca]\} \\
\text{ordered}_{ba}^{aa} &= \{[aa], [ba], [ab], [aa]\} \\
\text{ordered}_{aa}^{cc} &= \{[cc], [aa], [ac], [cc], [ca]\} \\
\text{ordered}_{ba}^{ab} &= \{[ab], [ba]\} \\
\text{ordered}_{ab}^{ca} &= \{[ca], [ab], [ba], [ac], [ca]\} \\
\text{ordered}_{ba}^{ca} &= \{[ca], [ba], [aa], [ac], [cb], [ca], [aa], [ab], [bc]\}
\end{aligned}$$

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

4.5.3 Rate of change of block types

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

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

$$\begin{aligned}
\frac{d\rho_{aa}^{aa}}{dt} &= R[ab \rightarrow aa] - R[aa \rightarrow aa] \\
\frac{d\rho_{aa}^{ca}}{dt} &= R[aa \rightarrow ca] + R[ba \rightarrow ca] + R[ab \rightarrow ca] \\
&\quad - R[ca \rightarrow ca] - R[ca \rightarrow cc] - R[ca \rightarrow ab] \\
\frac{d\rho_{ba}^{aa}}{dt} &= R[ca \rightarrow aa] + R[ab \rightarrow aa] + R[bb \rightarrow aa] \\
&\quad - R[aa \rightarrow aa] - R[aa \rightarrow ca] - R[aa \rightarrow ca] \\
\frac{d\rho_{aa}^{cc}}{dt} &= R[ca \rightarrow cc] + R[bc \rightarrow cc] - R[cc \rightarrow cc] - R[cc \rightarrow ab]
\end{aligned}$$

$$\begin{aligned}
\frac{d\rho_{ba}^{ab}}{dt} &= R_{ba}^{ca \rightarrow ab} + R_{bb}^{ab \rightarrow ab} - R_{ab}^{ab \rightarrow aa} - R_{ab}^{ab \rightarrow cb} \\
\frac{d\rho_{ab}^{ca}}{dt} &= R_{ab}^{aa \rightarrow ca} + R_{ca}^{ca \rightarrow ca} + R_{bc}^{ab \rightarrow ca} \\
&\quad - R_{ba}^{ca \rightarrow ab} - R_{ba}^{ca \rightarrow ca} - R_{ba}^{ca \rightarrow bc} \\
\frac{d\rho_{ba}^{ca}}{dt} &= R_{ab}^{aa \rightarrow ca} + R_{aa}^{cc \rightarrow ca} + R_{bc}^{ab \rightarrow ca} + R_{cb}^{ab \rightarrow ca} \\
&\quad - R_{ab}^{ca \rightarrow ca} - R_{ab}^{ca \rightarrow bc} - R_{ab}^{ca \rightarrow bc} - R_{ab}^{ca \rightarrow aa} \quad (4.11)
\end{aligned}$$

4.5.4 Block transformation rates


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

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

Example: $R_{rr}^{rr \rightarrow pr}$

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

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



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

original invasion pair.

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

$$R\left[\begin{smallmatrix} rr \\ rr \end{smallmatrix} \rightarrow \begin{smallmatrix} pr \\ rr \end{smallmatrix}\right] = 2r_p \rho_{pr} q_{r/rp}^{rr}.$$

Example: $R\left[\begin{smallmatrix} rr \\ sr \end{smallmatrix} \rightarrow \begin{smallmatrix} rr \\ rr \end{smallmatrix}\right]$

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

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

$$R\left[\begin{smallmatrix} rr \\ sr \end{smallmatrix} \rightarrow \begin{smallmatrix} rr \\ rr \end{smallmatrix}\right] = 2r_r \rho_{rs} (q_{r|rs}^{rr} + q_{r/sr}^{rr}).$$

Unique block transformations

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

$$\begin{aligned} R\left[\begin{smallmatrix} aa \\ aa \end{smallmatrix} \rightarrow \begin{smallmatrix} ca \\ aa \end{smallmatrix}\right] &= 2r_c \rho_{ca} q_{a/ac}^{aa} \\ R\left[\begin{smallmatrix} ca \\ aa \end{smallmatrix} \rightarrow \begin{smallmatrix} aa \\ ba \end{smallmatrix}\right] &= 2r_b \rho_{bc} q_{a/cb}^{aa} \\ R\left[\begin{smallmatrix} ca \\ aa \end{smallmatrix} \rightarrow \begin{smallmatrix} cc \\ aa \end{smallmatrix}\right] &= 2r_c \rho_{ca} (q_{a|ca}^{ca} + q_{c/ac}^{aa} + q_{a/ac}^{ca}) \\ R\left[\begin{smallmatrix} ca \\ aa \end{smallmatrix} \rightarrow \begin{smallmatrix} ca \\ ac \end{smallmatrix}\right] &= 2r_c \rho_{ca} q_{a/ac}^{ca} \\ R\left[\begin{smallmatrix} aa \\ ba \end{smallmatrix} \rightarrow \begin{smallmatrix} aa \\ aa \end{smallmatrix}\right] &= 2r_a \rho_{ab} (q_{a|ba}^{aa} + q_{a/ba}^{aa}) \\ R\left[\begin{smallmatrix} aa \\ ba \end{smallmatrix} \rightarrow \begin{smallmatrix} ca \\ ab \end{smallmatrix}\right] &= 2r_c \rho_{ca} q_{a/ab}^{aa} \\ R\left[\begin{smallmatrix} aa \\ ba \end{smallmatrix} \rightarrow \begin{smallmatrix} ca \\ ba \end{smallmatrix}\right] &= 2r_c \rho_{ca} (q_{a/ba}^{aa} + q_{a/ba}^{ca}) \\ R\left[\begin{smallmatrix} cc \\ aa \end{smallmatrix} \rightarrow \begin{smallmatrix} cc \\ ac \end{smallmatrix}\right] &= 2r_c \rho_{ca} (q_{c|ca}^{cc} + q_{c/ac}^{cc} + q_{c/ac}^{ca}) \\ R\left[\begin{smallmatrix} cc \\ aa \end{smallmatrix} \rightarrow \begin{smallmatrix} ca \\ ba \end{smallmatrix}\right] &= 2r_b \rho_{bc} (q_{c/cb}^{cc} + q_{c/cb}^{ca}) \\ R\left[\begin{smallmatrix} ab \\ ba \end{smallmatrix} \rightarrow \begin{smallmatrix} aa \\ ba \end{smallmatrix}\right] &= 2r_a \rho_{ab} (q_{b|ba}^{ab} + q_{b/ba}^{ba}) \end{aligned}$$

$$\begin{aligned}
R\left[\begin{smallmatrix} ab \\ ba \end{smallmatrix} \rightarrow \begin{smallmatrix} ab \\ bc \end{smallmatrix}\right] &= 2r_c \rho_{ca} q_{b/ac}^{ab} \\
R\left[\begin{smallmatrix} ca \\ ab \end{smallmatrix} \rightarrow \begin{smallmatrix} ca \\ aa \end{smallmatrix}\right] &= 2r_a \rho_{ab} (q_{c/a}^{ca} + q_{a/ba}^{ca}) \\
R\left[\begin{smallmatrix} ca \\ ab \end{smallmatrix} \rightarrow \begin{smallmatrix} ab \\ ba \end{smallmatrix}\right] &= 2r_b \rho_{bc} q_{a/cb}^{ba} \\
R\left[\begin{smallmatrix} ca \\ ab \end{smallmatrix} \rightarrow \begin{smallmatrix} bc \\ ac \end{smallmatrix}\right] &= 2r_c \rho_{ca} (q_{b/a}^{ca} + q_{c/ac}^{ab} + q_{b/ac}^{ac}) \\
R\left[\begin{smallmatrix} ca \\ ba \end{smallmatrix} \rightarrow \begin{smallmatrix} ca \\ aa \end{smallmatrix}\right] &= 2r_a \rho_{ab} (q_{c/b}^{ca} + q_{a/ba}^{ac} + q_{c/ba}^{aa}) \\
R\left[\begin{smallmatrix} ca \\ ba \end{smallmatrix} \rightarrow \begin{smallmatrix} aa \\ bb \end{smallmatrix}\right] &= 2r_b \rho_{bc} (q_{a/b}^{ca} + q_{b/cb}^{aa} + q_{a/cb}^{ab}) \\
R\left[\begin{smallmatrix} ca \\ ba \end{smallmatrix} \rightarrow \begin{smallmatrix} bc \\ ca \end{smallmatrix}\right] &= 2r_c \rho_{ca} (q_{b/ac}^{ca} + q_{a/ac}^{cb}) \\
R\left[\begin{smallmatrix} ca \\ ba \end{smallmatrix} \rightarrow \begin{smallmatrix} bc \\ ac \end{smallmatrix}\right] &= 2r_c \rho_{ca} (q_{b/a}^{ca} + q_{a/ac}^{bc} + q_{c/ac}^{ba})
\end{aligned} \tag{4.12}$$

4.5.5 The '2 × 2 block' approximation

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

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

$$q_{y|z}^{wx} = \frac{\rho_{yz}^{wx}}{\rho_z^x} = \frac{\rho_{yz}^{wx}}{\sum_{i,j} \rho_{jz}^{ix}}$$

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

$$q_{x'/yz}^{vw} = \frac{\rho_{xyz}^{vw*}}{\rho_{yz}^{vw}},$$

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

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

$$q_{x'/yz}^{vw} \approx q_{x|y}^{uv}.$$

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

$$q_{x'/yz}^{vw} = \frac{\rho_{xyz}^{vw*}}{\rho_{yz}^{vw}} \approx \frac{\rho_{xy}^{vw} \rho_{yz}^{w*}}{\rho_y^w \rho_{yz}^{vw}} = \frac{\rho_{xy}^{vw} \sum_i \rho_{yz}^{wi}}{\sum_{i,j} \rho_{yj}^{wi} \sum_{i,j} \rho_{yz}^{ij}}. \tag{4.13}$$




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

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

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

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

4.6 Behaviour of the 2×2 model

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

4.6.1 Stability

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

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

4.6.2 Ecosystem composition

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

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

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

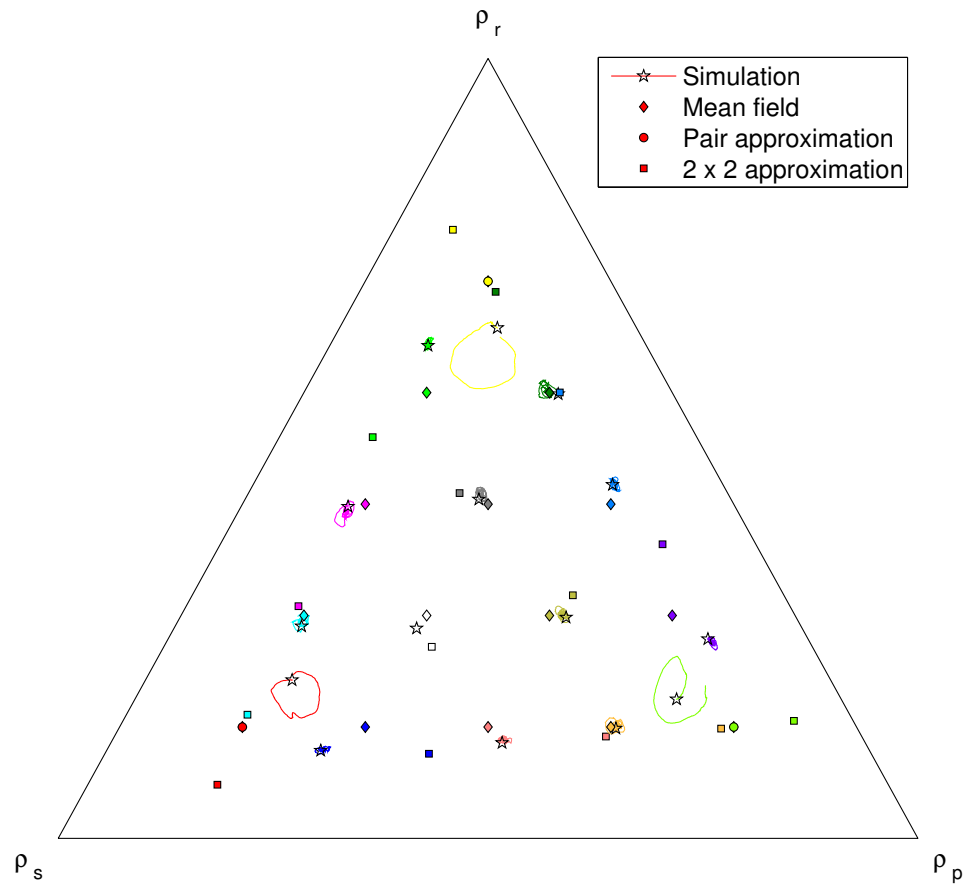


Figure 4.8: Comparison of fixed point locations of the species densities for 15 different values of the invasion rates under the mean field, pair and 2×2 approximation models, and SCA simulation. Each colour represents one set of invasion rates. The diamond, circle, and square symbols show fixed points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

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

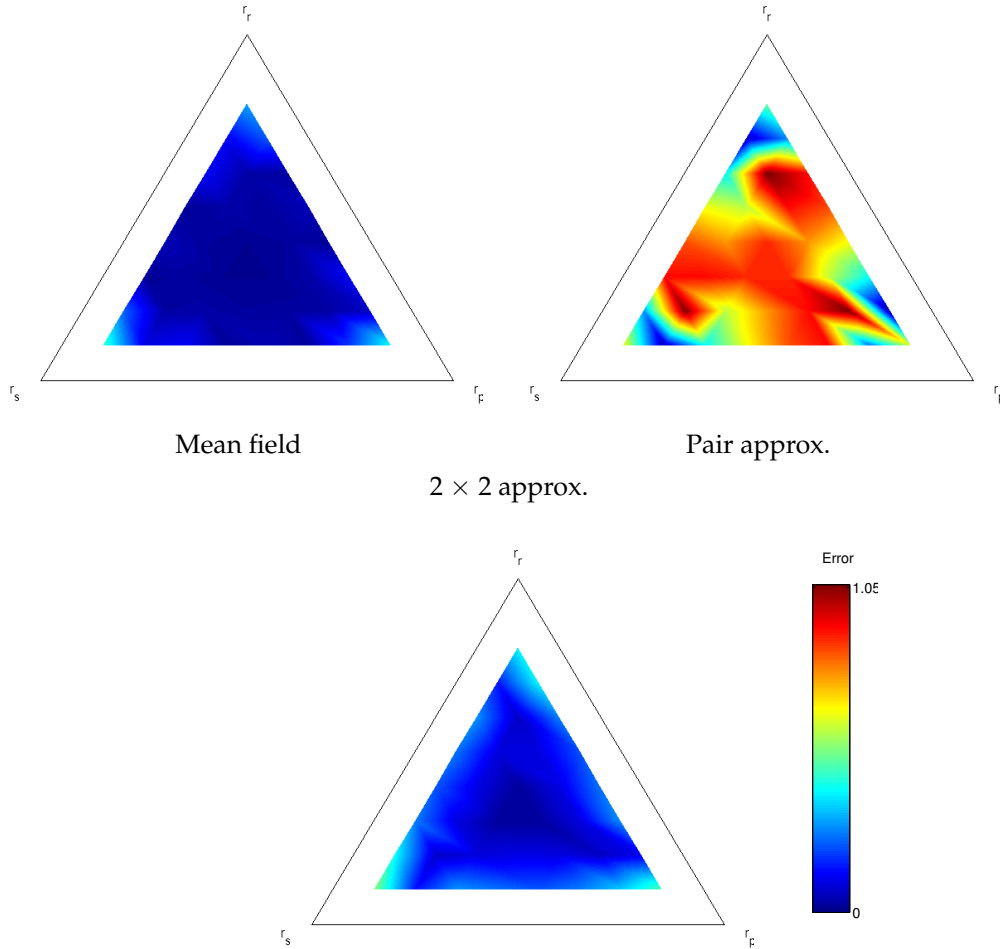


Figure 4.9: Difference in predicted species densities between each of the three approximations and the result of spatial SCA simulations on a 700×700 grid, for a range of invasion rates. The error measurement used in each case, and shown on the colour scale, is the distance in species-density space between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

4.6.3 Neighbour correlations and species clustering

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

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

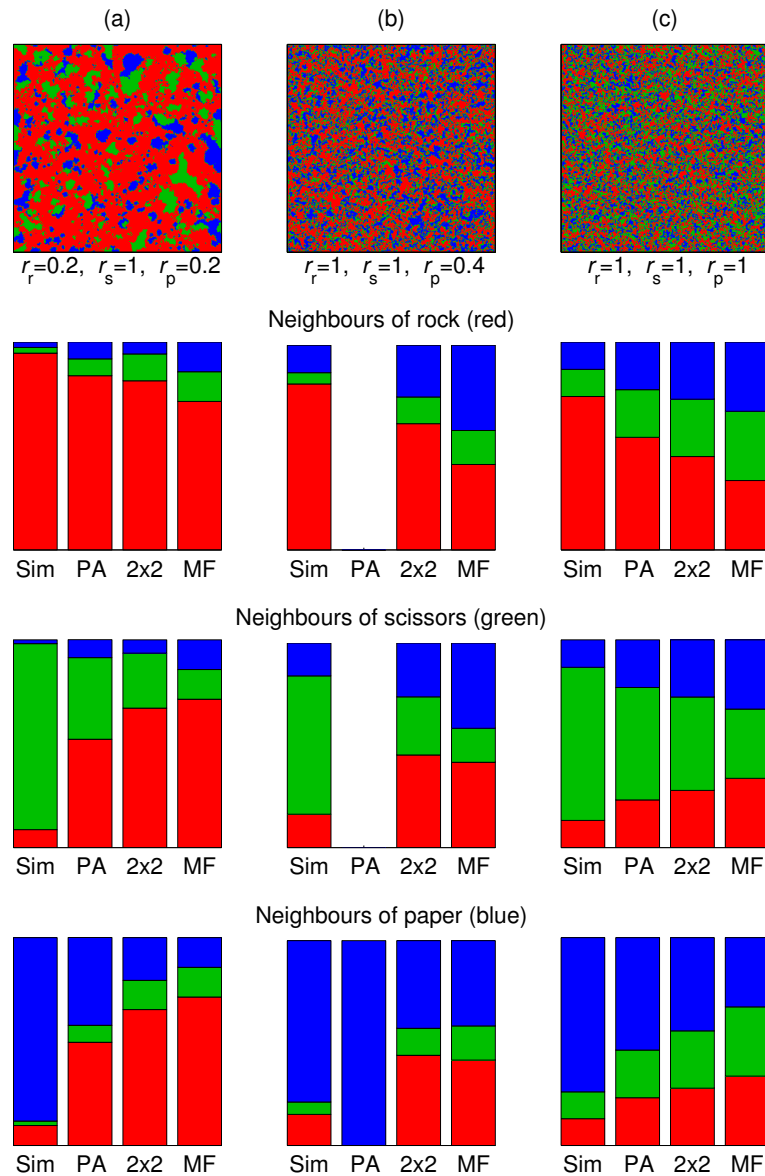


Figure 4.10: Predictions of densities of cell neighbours for three different sets of growth rates by simulation, 2×2 approximation, pair approximation and mean field techniques. The first row shows the results of explicit spatial simulations of the RPS ecosystem using three different sets of growth rates,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

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

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

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

Clustering and inaccuracy in densities

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

4.7 Discussion

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

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

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

Chapter 5

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

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

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

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

5.1 Explicit spatial stability

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

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

The spatial stability factor h

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

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

$$\begin{aligned}\dot{\rho}_r &= \rho_r(\rho_s r_r - \rho_p r_p)(1 - h) - \left(\rho_r - \frac{r_s}{r_r + r_s + r_p}\right)h \\ \dot{\rho}_s &= \rho_s(\rho_p r_s - \rho_r r_r)(1 - h) - \left(\rho_s - \frac{r_p}{r_r + r_s + r_p}\right)h \\ \dot{\rho}_p &= \rho_p(\rho_r r_p - \rho_s r_s)(1 - h) - \left(\rho_p - \frac{r_r}{r_r + r_s + r_p}\right)h\end{aligned}\quad (5.1)$$

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

5.2 Behaviour of the approximation

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

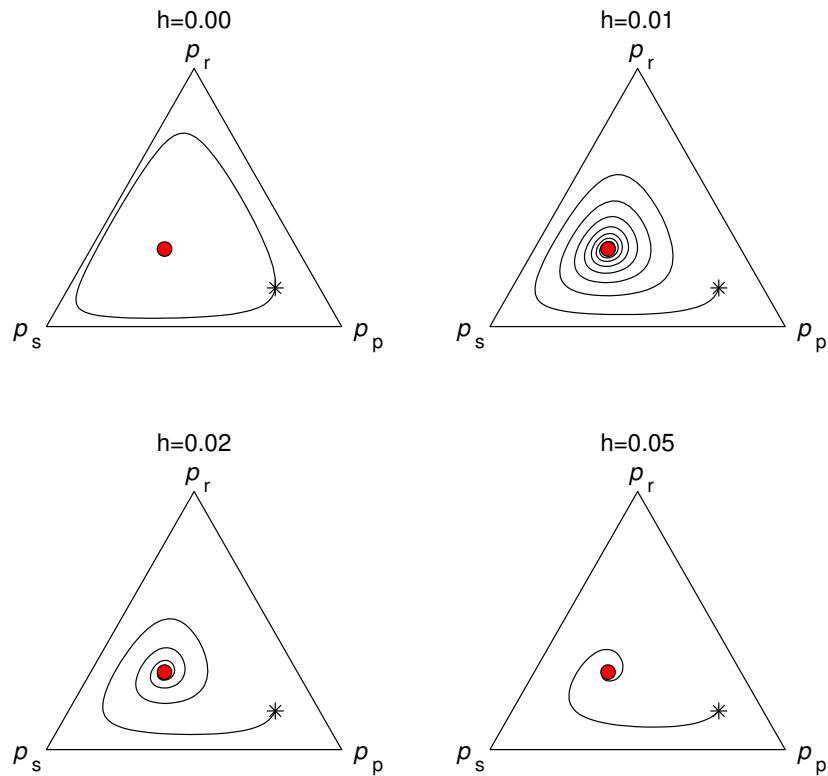


Figure 5.1: Behaviour of equations (5.1) for four values of h , with initial densities $(0.15, 0.15, 0.7)$ (marked by a star), and fixed invasion rates $(0.25, 0.3, 0.45)$, which determine the fixed point (red dot).

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

5.2.1 Extinction risk

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

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




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

5.2.2 Predictions of extinction risk

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

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

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

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

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

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

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

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

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