Simulating Invasions in a Spatially Structured Ecosystem

Carlo Maley
545 Technology Sq., NE43-803
Massachusetts Institute of Technology*
Cambridge, MA 02139, USA
email: cmaley@ai.mit.edu

Abstract

A recent distribution model[1] suggests that the success of ecological invasions diminish with the number of species in the resident community. However, such models generally lack any spatial structure or representation of the resources in the environment. The continuous distribution model of Case[1] was transformed into a configuration model that included a parameter for the dimensions of a spatially structured environment. This model allowed the comparison of the success of ecological invasions in a spatially non-structured (1 by 1) environment with environments where the organisms were distributed over a square grid. The addition of the explicit representation of resources in the model led to divergent results from Case's model. It was found that when the resource distribution in the environment was homogeneous, both spatial structure and community size reduced the probability of a successful invasion. However, the introduction of heterogeneity to the distribution of resources disrupted both of these effects.

1 Introduction

Case[1] developed a distribution[2] model of ecosystem invasions in a continuous state space. His analysis demonstrated that ecosystems of many species tended to repel invading species more often than ecosystems of relatively few species. This effect presumably falls out of the expectation that an ecosystem of many species will tend to more fully exploit the resources and thus leave fewer unutilized resources for invading species, as compared to an ecosystem with only a few species.

Case's analysis helps to highlight the importance of community level factors in the success of invasions. This contrasts with the more traditional focus on the properties of the invader itself. Of course, Case's model abstracts away most of the complexity of ecosystems and the dynamics of invasion. In particular, Case's model lacks spatial structure as well as the dynamics of the resources. Since invasions are spatial phenomena, we might ask how the Case's results would change if the ecosystems were situated in a spatially structured environment.

With this inspiration, I took Case's distribution model and simulated it in a configuration[2] model. Given such a simulation, we can begin relaxing some of the simplifying assumptions necessary for making the distribution model mathematically tractable. Specifically, I included both a representation of the resources and a "spatial dimension" parameter in the simulation so that ecosystems could be studied in varying degrees of spatially structured environments. I also allowed individuals to vary in their success in finding resources. Finally, the dynamics of invasion were investigated when the resources in the environment were heterogeneously distributed. We predicted that invasions would be more likely to succeed in a heterogeneous environment. Invaders should be able to locate, through dispersal of their offspring, nooks in the environment where the distribution of resources allow them to flourish.

2 Case's Model

Case's model is based on a matrix of Lotka-Volterra competition equations:[1]

$$\frac{dN_i}{dt} = \frac{r_i N_i}{K_i}(K_i - N_i - \sum_{j \neq i}^{n} \alpha_{ij} N_j) \tag{1}$$

Where there are n species all at one trophic level competing for m resources. He derives the α values by selecting random values from a log-normal distribution[3] for a species' utilization of each resource and then computing the amount of overlap from a method given by Levins[4]:

$$\alpha_{ij} = \frac{\sum_h p_{ih} p_{jh}}{\sum_h p_{ih}^2} \tag{2}$$

Where p_{ih} is the amount that species i utilizes resource h. All species in the model have intrinsic growth rates, $r = 0.5$. Case chose K_i such that there was an equilibrium for all species densities > 0. He then threw out any communities whose equilibriums were unstable.

*This research was supported in part by the generosity of the MIT AI Lab as well as DOD grant DAAH04-95-1-0557. I would like to thank W. Bossert, R. Brooks and M. Donoghue for their helpful advice.

[1] Case does not include the second N_i in his publication, although it is standard for the equation.

Given a stable equilibrium community of n species, Case introduced an invader species to the system in small numbers by adding a row and column to the competition matrix. He chose resource utilization values for the new species from the same distributions used to generate the resident community. An invading species was said to be successful if the community reached a new equilibrium where the invader species was still extant. Case found that the success of the invaders declined with the number of resident species in the community.

2.1 Criticisms

The Lotka-Volterra equations are oversimplified descriptions of competition. They attempt to translate an indirect interaction, through consumption of resources, into a direct interaction described by the α coefficients. According to Piank[5] "Indeed, the alphas in these equations may well be illusory and may often obscure the real *mechanisms* of competitive interactions."[emphasis his] Presumably, this was one reason Case took the trouble to describe the species by their utilizations of the different resources. To make the mathematical analysis tractable, and perhaps to locate the work within standard theoretical ecology, Case used these descriptions to derive the α values for the competition equations.

Unfortunately, important information is lost when a pair of species' resource utilization profiles are condensed down into a single α. This can be illustrated by considering two species utilizing two resources. Let species 1 require 10 units of resource 1 and 1 unit of resource 2. Similarly, let species 2 require 10 units of resource 1 and 9 units of resource 2. Assuming that each resource is replenished in equal amounts, resource 1 will become a bottleneck for the growth of the two species. And so a member of species 1 will have the same effect on the growth of species 2 as a member of species 2, since they both consume 10 units of the critical resource. In contrast, resource 2 will never restrict the growth of either species. So the actual formulation of the interaction between the two species should set $\alpha_{12} = \alpha_{21}$. However, equation 2 makes $\alpha_{12} \neq \alpha_{21}$ ($\alpha_{12} = 1.08$ and $\alpha_{21} = 0.60$) due to its inclusion of resource 2. Thus, a more accurate model of competitive interaction must take into account the dynamics of restrictive resources.

Adding more species to the community only complicates the dynamics, making some resources bottlenecks only under particular combinations of population levels. Population fluctuations in a community of competing organisms may stress different resources at different times. Tilman[7, 8] has taken these complexities to heart and developed a more elaborate model of competition and niche overlap.

Ricklefs and Schluter[6] criticize Case and others for the lack of spatial structure in their models. Because spatial structure introduces locality, what is treated as a homogenous population in the distribution model can become a heterogeneous collection of individuals, all situated within different networks of neighbors, in a configuration model. This variation, and the interactions between localities may well affect the dynamics of invasions.

3 The Configuration Model

Most of the criticisms of Case's model can be addressed in a spatially structured simulation of ecosystem invasions that includes a representation of the resources. I situated the population dynamics in a square grid of environment locations. Since the dimensions of the grid are parameters to the model, the degree to which the population is spatially structured can range from none (a 1 by 1 grid) to a highly distributed environment with only a few "organisms" in each environment location. In order to illuminate the effect of spatial structure on the dynamics of the model, the rest of the simulation was made to match Case's distribution model as closely as possible.

3.1 The Structure

An "organism" in the configuration model was described by its species membership and the amount of each kind of resource that it had gathered. A "species" was described by the amount of each kind of resource that was necessary for an organism to gather in order to survive and reproduce. The other important data structure in the model was an environment location. An environment location contained a set of organisms of the different species. It also contained two arrays of length m (the number of resources) describing the amount of each resource currently held in that environment location and also the amount of each resource that was added to this total every time step.

Each time step the environment was updated by adding a fixed amount to the resources in each environment patch. Then the organisms in each environment location were allowed to eat. If an organism failed to collected up the necessary amounts of the different resources, as described by its species, then that organism died by starvation. However, if an organism managed to collect up enough resources, it was allowed to reproduce by randomly placing a new organism of its species with no resources somewhere in its immediate neighborhood. The location for the offspring was selected by taking a 3 step random walk from the parent's position. The parent's resources were also reset to zero at this point.

3.2 The Eating Algorithm

When an organism was allowed to eat, it transferred resource units from its environment location into its internal store. Every organism tried to eat as much at it needed in a single time step, but when the resources ran out, no more organisms were allowed to eat during that time step.

This is the heart of the configuration model because it is the only path of interaction and competition between organisms. The particulars of the algorithm for exactly who got to eat in a time step had a strong effect on the dynamics of the model. Data were collected for two different consumption algorithms. Finally, data were collected with a heterogeneous distribution of resources in the environment.

3.2.1 Every Species Gets Its Share In order to eliminate any stochastic component to the model, and so better match the assumptions of Case's distribution model, one has to apportion out the resources to each species in proportion to the numbers of individuals in those species. However, given the fact that living and reproducing is an all or nothing affair, depending on the organism collecting all of its needs, one cannot simply apportion out the resources equally to all individuals. Instead, once the correct share was computed for a species, the individuals in that species were allowed to eat, one by one, until their share of the resources had been consumed.

3.2.2 Every Organism Gets a Chance A more realistic algorithm for resource collection is that individuals are randomly allowed to eat, regardless of their species, until all the resources are consumed. This introduces a large source of stochasticity and so makes the detection of community equilibrium non-trivial. However, equilibria are rarely observed in ecosystems. So this may not be much of a sacrifice.

3.2.3 Environmental Heterogeneity The random resource allocation to individuals was further elaborated by allowing the replenishment of resources in different patches to vary within 50% of the average. Thus, one environment location might have an abundance of resource 1 but little of resource 2, while the neighboring location might have an intermediate level of both resources. However, the replenishment levels for the resources were determined at the beginning of a trial and remained constant over time.

4 Methods

Data were collected on ecosystems with ten resources and either two, four, six, or eight resident species. Each of these community sizes was run under a 1 by 1 environment (i.e., no spatial structure), a 2 by 2 environment, a 4 by 4 environment, and an 8 by 8 environment. This was repeated for the "every species gets its share," "every organism gets a chance" and the "environmental heterogeneity" conditions.

At the beginning of a run, the needs of the species were generated randomly on a log-normal curve for each resource over the range from 0 to 10. A trial run started with four organisms (the same number as in an invasion) of each species in every environment location. The ecosystem was then run, for up to 1000 time steps, less if an equilibrium was reached. After this point, if any species had gone extinct, a new invader was introduced in that "slot." And in any case, at least one new invader was introduced. An invader was recorded as successful if the community reached equilibrium with the invader extant, or the invader was still extant after 300 generations. If a resident species went extinct during that period, it was replaced by a new invader. In this way, each invader experienced competition of a constant number of species, though there might have been more than one "invader" at any one time. Equilibrium was detected when all the species maintained their population levels within 30(generations). Once an invader was noted to be successful, it was eliminated from the population and replaced by a new invader. The same happened if the invader was unsuccessful and went extinct. This continued until a total of 10 invaders had been introduced into the community. The model was run between 20 and 70 times for each pair of community size and environment size, where each run consisted of 10 sequential invasions into the community.

An invasion began when four invaders were placed into a randomly chosen location. The amount of available resource was set so that across the entire environment a species could have a population of approximately 1000. However, a species that randomly had low resource needs would have a higher carrying capacity and a species that needed to consume a lot of resources would have a lower carrying capacity.

5 Results

A replication of Case's model was run in parallel with the "fair share" resource consumption algorithm in the non-spatially structured environment. Over a total of 4080 species in the resident communities, there was only a 0.31 correlation between Case's distribution model predictions and the actual population levels observed in the configuration model.

The probability for successful invasions for the four different environment sizes and four community sizes are summarized in Table 1 and depicted in Figures 1, 2, and 3.

"Success" was defined to be when the invader had a non-zero population level after the ecosystem reached an equilibrium, or in the cases where an equilibrium was not detected, when the invader was still extant after 300 generations. The data points in these graphs were generated by simply dividing the number of successful invasions under a pair of parameter settings by the total number of trials under those parameter settings.

Under the stochastic eating algorithm, the community size has a significant effect on the probability of success (p ¡ 0.005) but the apparent trend in spatial distribution is not significant at the $p = 0.05$ level. As might be inferred from Figure 3, neither community size nor spatial structure has a significant effect on the probability of success when the environmental resources are heterogeneously distributed.

6 Discussion

The three different resource conditions resulted in distinctly different effects. As Case indicated, increases in community size generally reduced the probability of successful invasions. However, the addition of environmental heterogeneity disrupted any clear effects in the model. Our prediction that environmental heterogeneity would increase the probability of invasion success were not borne out by the model.

6.1 Every Species Gets Its Share

The most striking effect under this condition is that the probability of successful invasion plummets when the en-

Competitors	*Fair Shares*	*Stoch. Eating*	*Hetero. Env.*
1 by 1 Environment			
2	0.430	0.105	0.045
4	0.313	0.051	0.088
6	0.338	0.040	0.091
8	0.278	0.056	0.060
2 by 2 Environment			
2	0.008	0.095	0.045
4	0.010	0.052	0.044
6	0.013	0.055	0.065
8	0.003	0.038	0.035
4 by 4 Environment			
2	0.008	0.065	0.075
4	0.000	0.067	0.055
6	0.000	0.050	0.045
8	0.000	0.050	0.055
8 by 8 Environment			
2	0.000	0.085	0.055
4	0.000	0.035	0.055
6	0.000	0.055	0.065
8	0.000	0.050	0.048

Table 1: These are the probabilities of successful invasion under three different conditions. The first column indicates the number of competitors that were in the community at the time of invasion. The second column gives the results when resources are consumed in proportion to the species population levels. The third column gives the results when organisms are selected randomly and are allowed to eat. And the last column gives the results when eating is stochastic but the resources have different abundances in the different environment locations.

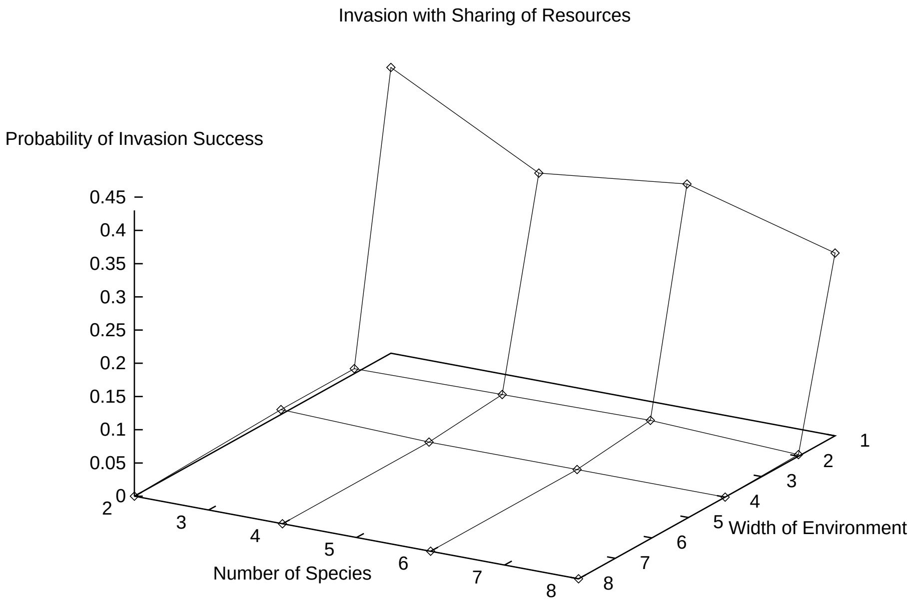

Figure 1: The percentages of trial runs where the invader was successful in establishing a colony at equilibrium as a function of the number of resident species and the number of patches along each side of a two-dimensional environment. When the width of the environment is 1, it is a spatially non-structured environment. That is, all organisms lived within a single environment location. In this condition, each species is allocated an amount of resources proportional to the number of individuals of that species.

Invasions Under Stochastic Resource Consumption

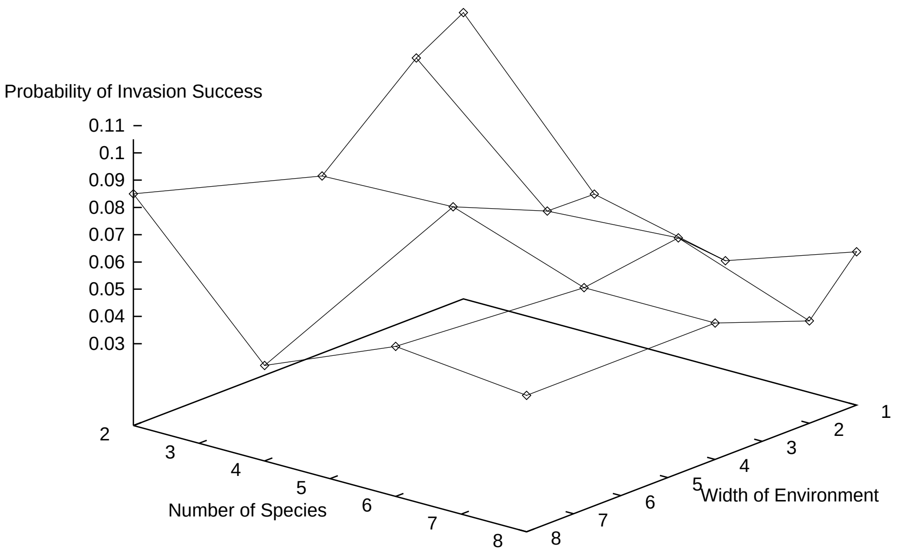

Figure 2: The consumption of resources was made stochastic by randomly choosing individuals in an environment location that had not yet eaten, and letting them consume resources. This was repeated until all the resources were consumed or all the organisms had eaten their fill.

Invasion Under a Heterogeneous Environment

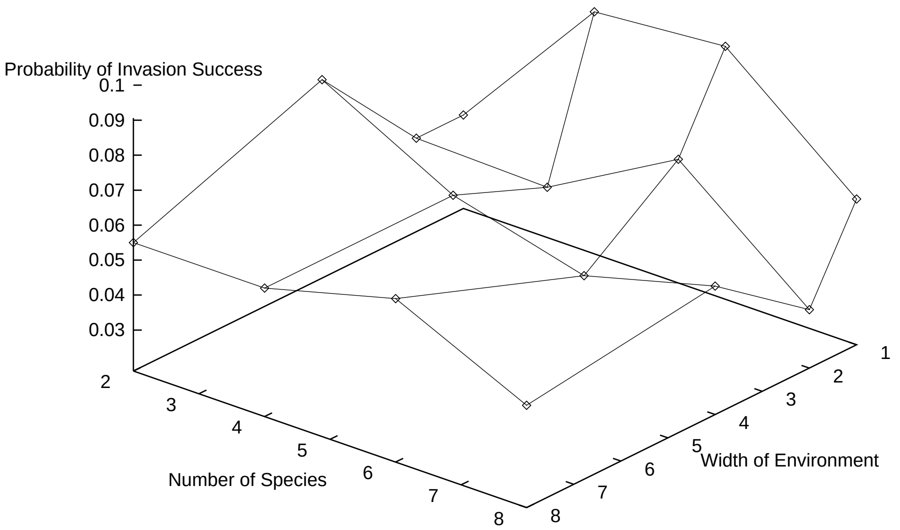

Figure 3: While the consumption of resources was still stochastic, the amounts of resources in an environment location varied across the environment.

vironment is given a spatial structure. This is due to the dispersion of the invaders across the environment. Although the invaders initially arrive in a cluster of four individuals, if any of those individuals reproduces, its offspring is likely to settle in a neighboring patch. With very small numbers of organisms, the invading species is allocated a tiny fraction of the resources. When this fraction is below the needs of a single organism of that species, the species goes extinct. In other words, this effect is mostly an artefact of the algorithm necessary to eliminate stochasticity in the resource consumption and thereby match Case's distribution model.

When there is no spatial structure, the probability of successful invasion decreases with the number of species in the community. This matches Case's results although his model predicts that the probability of successful invasion is more than twice the above results.

6.2 Every Organism Gets a Chance

When each organism in a location has an equal chance to consume resources, genetic drift becomes a dominant factor in the invasions, reducing the chance of success. Of course, this is true in nature as well, since invasions almost always involve small population sizes, and therefore genetic drift. In this case, the probability of successful invasion into a non-spatially structured environment (1 by 1) dropped by about a factor of five.

Figure 2 shows a general trend of greater success in the smaller communities with less spatial structure to less success when there were many competitors in a more spatially dispersed environment. It seems that both spatial structure and large community size inhibit invasions. The flatness of the 8 species community line hints that large community size mediates the effects of spatial structure on invasion success.

6.3 Environmental Heterogeneity

The introduction of spatial heterogeneity disrupted the effects of community size and spatial structure. It tended make the probability of successful invasion less sensitive to the number of competitors. However, it did not significantly raise the probability of success in most cases. The only exceptions were the 4 and 6 species communities in the non-spatially structured environments. In those two cases, the probability of success doubled. Since the non-spatially structured environment was not technically heterogeneous, it should probably be considered in isolation from the spatially structured environment results. In contrast, the probability of success in a 2 species community was about half that of the homogeneous environment, regardless of the dimensions of the environment.

6.4 Criticisms

If the interaction of spatial structure, community size, and environmental heterogeneity resulted in subtle effect, they would not be detected well given the variation in the data. A more comprehensive and sensitive investigation would be valuable.

The environment niches of the species, as described by the array of resources required for survival and reproduction, were inflexible. Species could not change their resource requirements, narrowing their niche, as a strategy for achieving coexistence in the presence of competitors with partially overlapping niches. In runs where niche flexibility was allowed, the invaders did indeed succeed in establishing a colony more often. Finally, reproduction and survival are intimately tied in the model, more for computational simplicity than for biological plausibility.

7 Conclusions

The inclusion of resources in the model created bottleneck dynamics that caused the actual population dynamics to diverge from the predictions of Case's model. When the resource distribution in the environment is homogeneous, both larger communities and greater "structure" of the environment tended to reduce the probability of successful invasions. However, environmental heterogeneity seems to interact strongly with community size and spatial structure such that they no longer have a clear effect on invasion success.

There are a number of future paths of exploration, even using this simple model. Most of these remain currently unexplored due to time constraints. The most obvious extension would be to push farther out into parameter space, and examine the dynamics of invasions into communities with more than eight species. The addition of sexual reproduction, and the subsequent low population density penalty may well change the dynamics of invasion significantly.

As Pianka[5] points out, there is a lot of work to be done on the mechanisms of competitive interactions. The implementation of resources and resource consumption is simplistic at best in this model. One fruitful course of investigation would be to make these dynamics more biologically reasonable in order to see how community structure and resource dynamics effect the phenomenon of invasion. It would then be interesting to investigate the dynamics of the species area curve over different sizes of environment.

In any case, the addition of spatial structure in a configuration model opens up a host of topics for exploration. Yet, as with all models, we can only take the results as suggestive and in need of verification in the field and laboratory.

References

[1] T. J. Case. Invasion resistance arises in strongly interacting species-rich model competition communities. *Proceedings of the National Academy of Science*, 87:9610–9614, 1990.

[2] Hal Caswell and A. M. John. From the individual to the population in demographic models. In D. DeAngelis and L. Gross, editors, *Individual-Based Models and Approaches in Ecology*, pages 36–61, New York, 1992. Chapman and Hill.

[3] D. A. S. Fraser. *Probability and Statistics*. Duxbury Press, North Scituate, MA, 1976.

[4] R. Levins. *Evolution in Changing Environments.* Princeton University Press, Princeton, NJ, 1968.

[5] E. R. Pianka. Competition and niche theory. In R. M. May, editor, *Theoretical Ecology*, pages 114–141. W. B. Saunders Co., Philadelpha, PA, 1976.

[6] R. E. Ricklefs and D. Schluter. Species diversity: Regional and historical influences. In R. E. Ricklefs and D. Schlute, editors, *Species Diversity in Ecological Communities*, pages 350–363. University of Chicago Press, Chicago, IL, 1993.

[7] D. Tilman. *Resource Competition and Community Structure.* Princeton University Press, Princeton, NJ, 1982.

[8] D. Tilman. *Dynamics and Structure of Plant Communities.* Princeton University Press, Princeton, NJ, 1988.