

Different Approaches to Theoretical Population Biology

Nir Friedman

12/14/08

Abstract

A brief discussion of the main ideas behind population biology. The basic starting points in modelling populations, and common assumptions are discussed. Possible difficulties in assumptions are also mentioned. Different models are discussed as examples, each using different (progressively less) assumptions. Finally, the possible dangers in some of these assumptions is discussed.

1 Introduction

The goal of population biology is to find ways of accurately modelling the populations of various living species in the wild. This can include not only the numbers of the species but also characteristics of the population such as genotype and phenotype. This problem is inherently attractive to practitioners of statistical mechanics because there are many similarities in the structure of the problem. In statistical mechanics the individual particles themselves are very complicated objects, just as individual organisms are. However, in population biology we are generally interested in statistical quantities that describe the population at large: the number of each of type of species, genetic make up, stability against changes in environment, and so on. These quantities are analogous to the macroscopic quantities of statistical mechanics. And just as in statistical mechanics we can hope to find good descriptions of these macroscopic quantities while only taking into account a small fraction of the complexity of individual particles, we can likewise hope to find accurate models of population biology without taking into account the hopeless complexity of individual organisms.

Population biology has a somewhat different flavour to it than the physics that the average physicist is used to. While making correct predictions is paramount in any field of science, in traditional physics we are generally guided by fundamental postulates. The consistency and correctness of these postulates, the mechanisms they lead to, and the challenges in extracting predictions about complicated systems from these postulates occupy much of traditional physics. In population biology the fundamental tenets behind models are incredibly rough approximations to begin with and do not pretend to represent anything fundamental about the organism in question (at least, nothing fundamental in a reductionist sort of approach). So one can ask: what has been learned when a correct (in the sense of predicting empirical data) model has been developed?

The answer is that in a reductionist sense nothing, i.e. we probably have not learned a thing about the individual organism, at least nothing that could not be much more easily verified in another way. However, the very fact that our model is correct gives us tremendous knowledge on what the important collective mechanisms are. Building a model of bee population from simple automatons on a grid will probably never teach you anything about the bee, however you one can learn what the dominant mechanisms are contributing to changes in the population distribution of the bees as a whole. One can see how changes in environment will impact them, how they will be affected by interactions with another species, and perhaps why they have evolved the way that they have. Doing this furthers our real understanding of the world around us, and perhaps in time could help to explain our own development as a species.

In addition, such models sufficiently far advanced would have plenty of practical applications. The ability for example to specifically and quantitatively find the factors responsible for the population decline of an endangered species might well allow us to more efficiently resuscitate their populations. Better understanding of the population cycles of bees and trout, animals that produce resources for humans, might allow us to better make use of these species to continue producing food.

There are two approximations that can be typically made in dealing with the modelling of a population. The first one is continuity. Rather than assuming discrete individuals, a continuous population is assumed. Secondly, the explicit spatial dependence of the problem is sometimes ignored. Making both of these assumptions clearly makes the problem much easier: instead of dealing with a multitude of individuals each with its own spatial coordinate as a function of time, we now just deal with a single number that is a function of time. However, both of these assumptions can lead to imprecisions as demonstrated by Durrett and Levi [2]. We turn now to look at examples of various models that result depending on the combinations of assumptions.

2 Mean Field Models

Mean field models are not only the simplest models of population dynamics, they are also the oldest. Since we have continuous variables with only time dependence, mean field models are about setting up systems of ordinary differential equations. One of the foundational papers in the field, due to Lotka, suggests the following system[1]. We take X_1 to be the total mass of some plant species in the system, and X_2 to be the total mass of some herbivorous species in the system. We then write:

$$\frac{dX_1}{dt} = A_1X_1 - B_1X_1X_2.$$

$$\frac{dX_2}{dt} = A_2X_1X_2 - B_2X_2.$$

In general, the coefficients A_i and B_i can depend on X_1 or X_2 , however we assume that the dependence is weak within a small range of values. This model takes the following assumptions to be approximately true (only approximately because we still allow the possibility of the coefficients weakly depending on the variables):

- In the absence of predators, plant growth is proportional to the existing amount of plant life
- In the absence of food, herbivores die off at a rate proportional to the number of herbivores
- The rate at which herbivores consume plants is proportional both to the number

This model is very simple and intuitive, yet still yields two interesting observations. The first one comes from examining the equilibrium point. We can see that there are two equilibrium points:

$$X_1 = X_2 = 0,$$

or the more interesting

$$X_1 = \frac{B_2}{A_2}, \quad X_2 = \frac{A_1}{B_1}.$$

Remember that the coefficients depend on the X 's, and so this is not necessarily a trivial problem to solve. It may have no solutions, or multiple solutions. However, suppose we have a solution somewhere. If we linearize around this equilibrium point, we will find a simple two dimension oscillation, with frequency given by $\omega = \sqrt{\bar{A}_1\bar{B}_2}$, where \bar{A}_1 and \bar{B}_2 are the equilibrium values of A_1 and B_2 [1]. This corresponds to the intuitive notion of populations of predator-prey systems oscillating around some mean value.

The second interesting observation comes from looking at the phase diagram. By dividing the two differential equations from before, we find:

$$\frac{dX_1}{dX_2} = \frac{X_1(A_1 - B_1X_2)}{X_2(A_2X_1 - B_2)}.$$

Let us picture X_1 on the y-axis and X_2 on the x-axis. Suppose for a moment that the herbivore population is increasing and the plant population is getting dangerously low. On our graph, this means we are moving down and to the right (negative slope), somewhere slightly above the positive y-axis. As X_1 goes to zero however, so does the slope of the line. Hence the phase curve can never actually touch the y-axis. It will not even be asymptotic. As X_1 becomes very small, the slope gets very close to zero and hence X_2 is free to continue increasing without much impacting X_1 . However once X_1 is small enough and X_2 is large enough, in the expression for the slope the numerator will be negative and the denominator will be negative as well. Hence the slope will turn positive at some point and the curve will turn away from the positive y-axis. Thus we conclude that X_1 and X_2 will always be strictly positive (provided they start that way)[1]. This means that this model does not allow for the possibility of extinction.

3 Mean Field Models with Diffusion

In models such as these, we remove the assumption of spatial independence that we made earlier. We now have a continuous variable that represent population density as a function of space and time. Hence our problems are specified by partial differential equations. Here we consider a problem taken up by Dahmen et al [3]. We consider normal bacteria growth, however with a growth constant that depends on space as well as a “wind” flow that causes the bacteria to drift. We have the equation

$$\frac{\partial}{\partial t}c(\mathbf{x}, t) = D\nabla^2c(\mathbf{x}, t) - \mathbf{v} \cdot \nabla c(\mathbf{x}, t) + U(\mathbf{x})c(\mathbf{x}, t) - bc^2(\mathbf{x}, t).$$

On the right hand side, the term involving U is the familiar term of growth rate which is proportional to the amount of bacteria already present. However here the growth rate depends on position explicitly. The

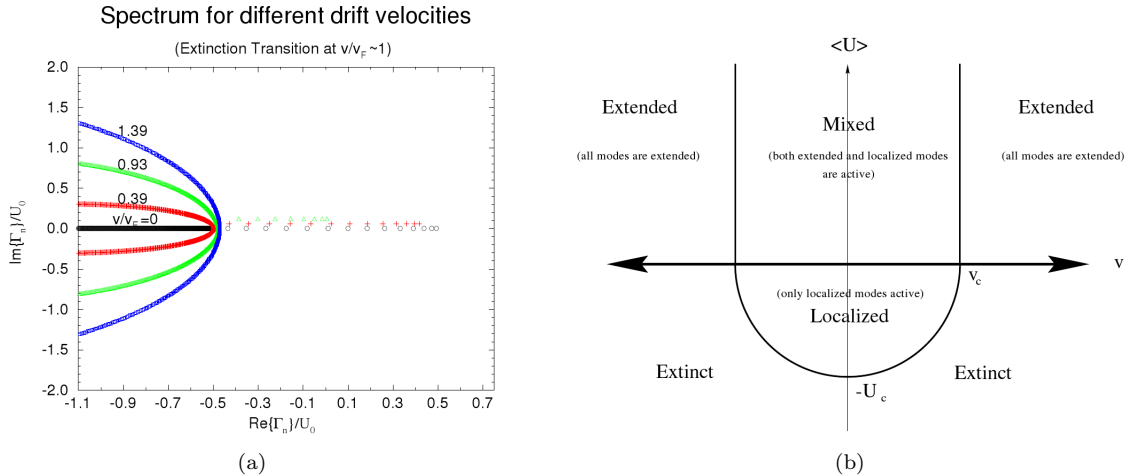


Figure 1: **(a)** The complex eigenvalues of the linear operator \mathcal{L} . A value of $U = 0.5$ was taken in the oasis and $U = -0.5$ in the desert. Our dimensional reference velocity is $v_F = 2\sqrt{U_0 D}$. These simulations used a grid with 1000 sides with an oasis of width 20. Hence the average value of U is $\langle U \rangle = -0.48$. Note that this is the value where the right end of all the parabolas intersect (the mobility edge): in other words the largest real value of the continuous spectrum is equal to $\langle U \rangle$. **(b)** A schematic phase diagram. If $\langle U \rangle < 0$ then there cannot be a surviving extended phase because the mobility edge is negative. On the other hand for large values of v , there is no discrete spectra (as we saw from (a)) and hence there is no localized state [3].

Laplacian term shows the tendency of the bacteria to diffuse. The term involving \mathbf{v} , which is the velocity of the flow, represents the effect of drift. Finally the last term is a term which prevents excessive crowding of the bacteria into one place.

We further deal the idea of an ‘‘oasis’’: a small region in which U is positive, and outside this region U is negative. In the long term, the bacteria can only survive in the region where U is positive, however they have trouble staying in the oasis because the wind keeps blowing them out. We are interested here in the idea of a phase transition taking place: we can imagine that if \mathbf{v} is very large, the bacteria will get blown out of the oasis too quickly and will be unable to survive. On the other hand if \mathbf{v} is zero then the bacteria will clearly thrive within the oasis and stay alive. Hence somewhere in the middle a transition must take place.

Since we are investigating the transition near extinction, we can linearize our equation by assuming that c is small. hence we have

$$\frac{\partial}{\partial t} c(\mathbf{x}, t) = D \nabla^2 c(\mathbf{x}, t) - \mathbf{v} \cdot \nabla c(\mathbf{x}, t) + U(\mathbf{x}) c(\mathbf{x}, t) = \mathcal{L} [c(\mathbf{x}, t)],$$

where \mathcal{L} is the linearized operator of interest. If we write $c = u(\mathbf{x})e^{\lambda t}$, then we have the eigenvalue problem

$$\lambda u(\mathbf{x}) = \mathcal{L} u(\mathbf{x}).$$

Notice how similar this equation is to the time independent Schroedinger equation. In this case however the analysis is greatly complicated by the presence of the term associated with \mathbf{v} ; this cause \mathcal{L} to be non-Hermitian. However we can still diagonalize it and find left and right eigenvectors as it turns out[3]. Figure 1(a) presents the results of numerically finding the spectrum. As we can see, in general there is a discrete spectrum (which corresponds to localized states around the oasis) and a continuous spectrum (which corresponds to delocalized states; note that we are using periodic boundary conditions). There are several important features on this diagram.

First of all, we see that for very large velocities the localized states vanish completely. Secondly, the largest real value of the continue spectrum does not depend on v very much. In fact this quantity (referred to as the mobility edge) is equal to the average value of U over the space in question. Finally, any value of




Figure 2: **(a)** On the vertical axis, $\overline{t_{ex}}$ is the extinction time. As we can see, larger mutation rates have a negative effect on survivability in constant environment. Vertical lines are the estimated phase transitions between the alive and extinct phases. **(b)** We can see the distribution in the trait over the population for various time intervals and mutation amplitudes. Note that the thick line with $M = 0.1$ is a steady state and hence doesn't have an associated time interval. With mutation rate zero, notice how the optimum trait is converged on while positive mutation rates lead to a wider distribution. [5]

v for which there are no positive eigenvalues is a phase that leads to extinction. Thus we can make a phase diagram using v and $\langle U \rangle$ as our parameters, shown in Figure 1(b).

This model has been tested using the species of bacteria *E. coli* RW 120 by Lin et al. The experimental set up used the bacteria in a very long thin tube, on which was shone light with UV in its spectrum. In addition a UV light filter was put in front of the light and moved back and forth. Because UV light is deadly to this bacteria species, the area covered by the filter constituted the “oasis”. Note also that moving the oasis is precisely equivalent to having a drift term moving the bacteria, as can be seen by a Galilean transformation. In addition to using experimental methods, a more thorough numerical simulation was done. The qualitative features of the model due to Dahmen were verified, however it was found that in the full non-linear analysis that there was no mixed state, i.e. that only extended modes would be active in that portion of the phase diagram. Another interesting feature not captured by any of the models was the appearance of double peaked and oscillating distributions of bacteria within the localized state[4]. However overall the the experiment showed the potential of these types of model and verified in this case the validity of taking the population to be a continuous variable (which is not surprising, given the massive numbers of bacteria populations).

4 Individual Based Models

In the previous example, we were dealing with a population that was so large (bacteria cultures) that it was very reasonable to ignore the effects of individuals. Here we examine a model which is right now primarily theoretical. For a variety of reasons it is harder to gather reliable biological data about non-microbial animals: smaller populations mean larger noise, longer breeding time, larger physical size makes isolated environment impossible, etc. Hence this theory has not yet found experimental comparisons, however hopefully in the future relevant data will be found. This model also has some complexities beyond keeping track of individuals. Rather than simply having a growth favorability U as function of space and time, it has an environment (a function of time only) that individuals can adapt to. Hence it models the genetic characteristics of the population as well.

The model takes place in two dimensions, on an $N \times N$ grid. Each spot on the grid can be occupied by at most one individual. Each individual has a “trait” characterizing him; the i th individual has trait z_i . This number is fixed at birth and is bounded on the unit interval. The changing environment is denoted by the variable $\phi \in [0, 1]$. The fitness of an individual, $f_i = 1 - |z_i - \phi|$, is a measure of how well adapted the individual is to its environment[5].

At each step in time, each individual gets to “move”. Moving consists of several steps:




Figure 3: As we can see in this case, very large values of \mathcal{M} still give very poor results however now the best survivability is actually obtained for $\mathcal{M} = 0.1$.

1. A check to see if the individual survives. The probability of individual i dying is given by

$$p_i = p_0 \left[1 - \exp\left(-\frac{\mathcal{S}}{f_i}\right) \right].$$

Both p_0 and \mathcal{S} are control parameters: p_0 is the maximum probability of extinction which will be the case for an individual with $f_i = 0$, and \mathcal{S} is called the selection pressure of the environment. It controls both how severely individuals are punished for straying from optimal fitness as well as the maximum probability of survival. Note that changes in p_0 are equivalent to changes in the time interval of a step.

2. An actual movement. That is the individual will randomly move to any of his nearest neighbours on the lattice, provided they are empty. If none are empty, no movement occurs.
3. If movement brings the individual in proximity with one more other individuals, then one is chosen at random for mating. The result of mating is the populating of the grid with new individuals. The number of children is equal to either the number of empty nearest neighbours to both parents (note that there are at most 6 such squares), or to N_{off} , whichever is smaller.
4. The children receive the traits as a combination of their parents genes and a random mutation factor. If k is the child of i and j , we have

$$z_k = \frac{1}{2} (z_i + z_j) + m_k,$$

where m_k is a uniform random number distributed on $[-\mathcal{M}, \mathcal{M}]$. This variable is known as the mutation amplitude. Note that if z_k is generated outside of $[0, 1]$ it is brought back to the nearest of those two endpoints.

The system is specified by the various parameters mentioned throughout (which are all constants), with the exception of ϕ which as mentioned before can have time dependence. What is interesting to note is that when one attempts to build a mean field equivalent of this model, it is possible to do so up to the point of dealing with traits. There is no way within mean field theory to predict what the expected trait value will be, and hence the only way to proceed with creating the mean field theory is to take a single fixed z as a constant [3]. Thus the role of the mutation amplitude is eliminated completely from the model. So the individual based model takes into account the mutation rate, even though mean field approaches are completely unable to do so. The results between a mutation rate of zero and the mean field approaches are quite similar in some circumstances, given the correct initial conditions.

The first set of simulations regarded cases where ϕ is held constant (at 0.5), we can see the impact of raising the mutation rate in figure 2. Note that the initial population has random initial traits. We can see however that in all cases a lower mutation rate is better. When the environment is constant, every generation is better adapted than the previous one (survival of the fittest). However mutations limit how

well adapted a population can be; at a certain point well adapted parents are having children that are more poorly adapted than they are due to random fluctuations.

The second set of simulations were for ϕ being some periodic function. In this case it was found that there was an optimal value of the mutation amplitude to give the best chance of survival (i.e. optimize the range of values for \mathcal{S} over which the population would survive). Intuitively, the mutation amplitude needs to be large enough to help adapt to the new amplitude faster, but not so large that it destroys the ability of the population to adapt to a specific trait. The larger the frequency of oscillation of ϕ , the larger the appropriate value of \mathcal{M} . Figure 3 shows the equivalent to figure 2(a) in the case of an oscillating environment.

5 Comparison of Models

In all these models, although we justify the approximations heuristically, we do not rigorously justify the approximations being made. Durrett et al undertook a very careful examination into the different kinds of models and found surprising cases where the models differ. They consider a spatial version of one of Maynard Smith's evolutionary games. In this case they consider the Hawk-Dove game. Suppose that a hawk receives a payout A_{HH} when interacting with another hawk, and a payout A_{HD} when interacting with a dove. The formalism is similar for doves. Now assume that there is an inherent death rate in the system proportional to the density of individuals, and a further range of growth or death that is proportional to the interaction in the game. In other words, in the mean field theory version we have

$$\frac{dH}{dt} = H \left\{ a \frac{H}{H+D} + b \frac{D}{H+D} - \kappa(H+D) \right\},$$

with similar equations for doves. The first two terms describe the average payout for interacting with another creature: the utility times the probability is summed over the two cases. The last term is just the death rate.

Using analogous assumptions, 3 other types of models are constructed: a mean field theory with diffusion terms, a "patch" theory that involves individuals residing in different areas but with no explicit spatial connection between the patches, and an individual based model. It is found that for this exact theory, different theories may give either the same or different results depends only on the choice of the payoffs, A_{ij} [2].

6 Conclusion

Population biology provides the opportunity to do some interesting modelling by trying to take reasonable first principles and then calculating system dynamics. In some cases we can find interesting behaviours reminiscent of statistical mechanics, such as phase transitions, stable and unstable equilibria, and so on. Hopefully in the future with further comparison to experimental results, more accurate predictions will be possible.

References

- [1] A. J. Lotka, *Proc. Nat. Acad. Sci. USA* **6**, 410 (1920).
- [2] Richard Durrett and Simon Levi, *Theoretical Population Biology* **46**, 363-394 (1994).
- [3] Karin Dahmen, David Nelson and Nadav Shnerb, *Journal of Mathematical Biology* **41**, 1-23 (2000).
- [4] Anna L. Lin, Bernward A. Mann, Gelsy Torres-Oviedo, Bryan Lincoln, Josef Kas, and Harry L. Swinney, *Biophysics Journal* **87**, 75-80 (2004).
- [5] Ioana Bena, Michael Droz, Janusz Szwabinski, and Andrzej Pekalski, *Physical Review E* **76**, 011908 (2007).