



1-2008

The Robust Volterra Principle

Michael Weisberg

University of Pennsylvania, weisberg@phil.upenn.edu

Kenneth Reisman

University of Pennsylvania

Follow this and additional works at: http://repository.upenn.edu/philosophy_papers

 Part of the [Philosophy Commons](#)

Recommended Citation

Weisberg, Michael and Reisman, Kenneth, "The Robust Volterra Principle" (2008). *Departmental Papers (Philosophy)*. 5.
http://repository.upenn.edu/philosophy_papers/5

This paper is posted at ScholarlyCommons. http://repository.upenn.edu/philosophy_papers/5
For more information, please contact repository@pobox.upenn.edu.

The Robust Volterra Principle

Abstract

Theorizing in ecology and evolution often proceeds via the construction of multiple idealized models. To determine whether a theoretical result actually depends on core features of the models and is not an artifact of simplifying assumptions, theorists have developed the technique of robustness analysis, the examination of multiple models looking for common predictions. A striking example of robustness analysis in ecology is the discovery of the Volterra Principle, which describes the effect of general biocides in predator-prey systems. This paper details the discovery of the Volterra Principle and the demonstration of its robustness. It considers the classical ecology literature on robustness and introduces two individual-based models of predation, which are used to further analyze the Volterra Principle. The paper also introduces a distinction between *parameter robustness*, *structural robustness*, and *representational robustness*, and demonstrates that the Volterra Principle exhibits all three kinds of robustness.

Disciplines

Philosophy

The Robust Volterra Principle*

Michael Weisberg and Kenneth Reisman†‡

Theorizing in ecology and evolution often proceeds via the construction of multiple idealized models. To determine whether a theoretical result actually depends on core features of the models and is not an artifact of simplifying assumptions, theorists have developed the technique of robustness analysis, the examination of multiple models looking for common predictions. A striking example of robustness analysis in ecology is the discovery of the Volterra Principle, which describes the effect of general biocides in predator-prey systems. This paper details the discovery of the Volterra Principle and the demonstration of its robustness. It considers the classical ecology literature on robustness and introduces two individual-based models of predation, which are used to further analyze the Volterra Principle. The paper also introduces a distinction between *parameter robustness*, *structural robustness*, and *representational robustness*, and demonstrates that the Volterra Principle exhibits all three kinds of robustness.

1. Introduction. Complex biological phenomena rarely admit of single, fully unified, theoretical treatments. For various reasons, theorists often study biological systems by investigating a family of different but related mathematical models. One motivation for investigating a family of related models is that it is a strategy for coping with highly idealized scientific theory. Biological theory routinely incorporates simplifying assumptions, such as the assumptions that a population is infinite in size or has non-overlapping generations. Such assumptions are convenient ways to deal with limited computational resources, and to make theory compact and

*Received September 2006; revised May 2007.

†To contact the authors, please write to: Michael Weisberg, Department of Philosophy, University of Pennsylvania, 433 Logan Hall, Philadelphia, PA 19104; e-mail: weisberg@phil.upenn.edu; Kenneth Reisman, Pluribo, Inc., 100 Park Avenue, Suite 1600, New York, NY 10017; e-mail: ken@pluribo.com.

‡Earlier versions of this paper were presented at the Australasian Association of Philosophy, the London School of Economics, and the University of Bristol. The authors wish to thank those audiences as well as Patrick Forber, Ken Waters, Deena Skolnick Weisberg, Uri Wilensky, and Bill Wimsatt for many helpful comments. Special thanks to Giacomo Sillari for his assistance in translating Volterra's original paper and his insightful thoughts about Volterra's aims and methods. Some of the research in this paper was supported by NSF grant SES-0620887 to MW.

Philosophy of Science 75 (January 2008) pp. 106–131. 0031-8248/2008/7501-0006\$10.00
Copyright 2008 by the Philosophy of Science Association. All rights reserved.

intelligible. The drawback of idealized models is that theorists must confront the problem of determining “whether a result depends on the essentials of the model or on the details of the simplifying assumptions” (Levins 1966, 20). Are the results generated by a model reliable, or are they artifacts of the analysis?

Another motivation for investigating a family of related models is that it is a strategy for understanding the generality of a result. Biologists often value results that are general—for example, a theoretical treatment of a system that remains true under many possible states of the system, or a result that applies to a wide range of different systems. Recognizing that any body of theory will depend on some set of assumptions, biologists possessing a general result will often want to know whether it will continue to apply under differing assumptions about the system.

To address these issues raised by the continuing investigation of multiple models, theorists have developed the technique of *robustness analysis*. This technique involves studying a number of similar, but distinct models of the same phenomenon, trying to find common predictions among them. These models may highlight different causal features, be formulated at different levels of abstraction, or even employ different mathematical frameworks in their representations of biological systems. In his famous discussion of robustness analysis, Levins describes what happens when the same prediction is made using multiple models.

[I]f these models, despite their different assumptions, lead to similar results, we have what we can call a robust theorem that is relatively free of the details of the model. Hence, our truth is at the intersection of independent lies. (Levins 1966, 20)

There is a small but growing philosophical literature about robustness analysis. Starting from the pioneering work of Levins (1966) and Wimsatt (1981), recent discussions have clarified the aims and methods of robustness analysis, and discussed whether robustness analysis has a role in confirmation, and exactly what this confirmation-theoretic role consists of.¹

1. Despite several authors addressing the confirmation-theoretic role of robustness, there has been little consensus among them. Orzack and Sober (1993) have argued that robustness analysis can play no non-redundant role in confirmation. Odenbaugh (2007) argues that robustness analysis allows us to discharge idealizations, showing that the model would make similar predictions if more realistic assumptions had been included. Weisberg (2006) argued that robustness analysis plays a role in confirmation, but the discovery that a theorem is robust is not a form of non-empirical confirmation. Rather, robustness analysis allows theorists to isolate particular properties or behaviors that will be present whenever a particular causal structure is instantiated. Forber (2007) takes a different tack, arguing that robustness plays a role in culling down possibilities

Despite having a schematic for robustness analysis and despite the important discussions of the possible confirmation-theoretic role of robustness analysis already in the literature, the philosophical literature has few if any detailed discussions of the actual robustness analyses carried out by scientists. Thus the primary purpose of this paper is to explain in detail the analysis and justification of an important biological principle called the *Volterra Principle* by robustness analysis. We will discuss how the Volterra Principle was discovered, why ecologists believe it to be true, and conduct some further robustness analysis by introducing two novel, individual-based models of predation. In addition, we will introduce a distinction between *parameter robustness*, *structural robustness*, and *representational robustness*, and show that the Volterra Principle exhibits all three kinds of robustness. Taken together, these three kinds of robustness analysis are a powerful way of demonstrating that a particular modeling result is not dependent on the particular assumptions or idealization embodied in a model or family of models.

2. The Lotka-Volterra Model of Predation. Predation is a much studied ecological phenomenon.² It is of great interest to ecologists because it often represents a force that keeps populations below their environment's carrying capacities. It is also a factor that can account for oscillation and other periodic dynamics of populations in which there is no external stimulation, such as in unchanging environments (Ricklefs and Miller 2000). Theoretical ecologists are interested in studying how predation leads to these phenomena. They construct models to study those factors that control the maximum population size as well as the phase, amplitude, and frequency of oscillations in populations. Naturally, some of these factors must be determined empirically, but there is also much that can be learned by analyzing clusters of models.

We have chosen to focus our discussion on predation because it provides an especially striking example of a robust theorem called the Volterra Principle. This principle was discovered by Vito Volterra, one of the found-

before empirical testing takes place. Our discussion is not primarily about the confirmation-theoretic role of robustness, although we will comment on the issue in various places. While our view of robustness analysis' confirmation-theoretic status is closest to the one articulated in Weisberg 2006, we believe that this analysis is only part of the story. The focus of this paper is the role of robustness analysis in showing which modeling assumptions are central and which are irrelevant for the production of a modeling result.

2. For a comprehensive review of the classical literature, see Royama 1971. For more contemporary discussions, including the history of predator-prey modeling, see Berryman 1992, Donalson and Nisbet 1999, Hanski et al. 2001, and Briggs and Hoopes 2004.

ers of mathematical biology, and has been subsequently discussed by many key figures in contemporary mathematical ecology including MacArthur and Connell (1966), Maynard Smith (1974), Roughgarden (1979, 1997), and May (2001). These theorists do not always use Levins' term 'robust theorem', but their discussions of predation, biological control, and the Volterra Principle make it clear that they believe the principle is robust. Before investigating this principle, we begin by discussing the model from which the principle was first discovered.

Volterra (1926a, 1926b) and Lotka (1956) independently proposed the first model of predator-prey interactions that we will discuss. This is probably the simplest possible model of predator-prey interactions, but even this simple model already displays rich dynamics as well as the property of greatest interest to us in this project. Volterra was explicit about the grounds for constructing such a simple model. He wrote:

As in any other analogous problem, it is convenient, in order to apply calculus, to start by taking in to account hypotheses which, although deviating from reality, give an approximate image of it. Although, at least in the beginning, the representation is very rough . . . it is possible to verify, quantitatively or, possibly, qualitatively, whether the results found match the actual statistics, and it is therefore possible to check the correctness of the initial hypothesis, at the same time paving the way for further results. Hence, it is useful, to ease the application of calculus, to schematize the phenomenon by isolating those actions that we intend to examine, supposing that they take place alone, and by neglecting other actions. (1926b, trans. G. Sillari)

To understand these remarks and the ways that the Lotka-Volterra model is a very simple way of representing predation, it is useful to think along the lines of a modeler approaching the problem for the first time. We ask: What are the essential quantities and interactions that our model needs to keep track of in order to represent predation?

If we are going to treat predation as a population-level phenomenon, as Lotka and Volterra did, then the primary quantities to keep track of are the size of the predator and prey populations. Alternatively, we can keep track of the population density, a quantity more easily measured empirically. We will refer generically to these measures as *species abundance*.

The next step in thinking about the structure of the model is to describe the intrinsic population dynamics of each species, or how the abundance of each species changes over time. Because the two species interact, their population dynamics are coupled together in the following way: The predators decrease the population of prey by eating them, while the prey

increase the population of predators by providing food. Abstractly, the relationship is one of negative feedback. Predators are negatively coupled to the prey, but prey is positively coupled to the predators (Maynard Smith 1974).

In order to construct simple, population-level, models of predation, we have six things to keep track of: the predator growth and death rates, the prey growth and death rates, the effect of predation on the population of prey, and the effect of prey capture on the population of predators. If we set up our model in terms of rates of increase and decrease, we can collapse intrinsic growth and death rates in to a single growth rate for the prey and, a bit less realistically, a single death rate for the predators. This will give us four quantities to keep track of.

Let V stand for the size of the prey population and P for the size of the predator population. If we express these basic relationships with coupled differential equations then we get the following basic equations (after Roughgarden 1997):

$$\frac{dV}{dt} = [\text{prey growth rate}] - [\text{capture rate of prey per predator}]P \quad (1)$$

$$\frac{dP}{dt} = [\text{predator births per capture}]P - [\text{predator death rate}]. \quad (2)$$

These equations provide a template for a large but tightly linked family of models. Starting from the simple possibilities, the prey growth rate could be linear, exponential, geometric, or logistic. The most typical death rate of the predators in predation models is constant, implying an exponential decay in the absence of prey. More complicated rate expressions are also possible, including functional dependence on environmental parameters and logistic decay when multiple sources of food are present. For the sake of simplicity, we will only consider examples where the predator death rate is constant, but we can modify the form of the intrinsic prey population growth rate.

Of greater biological interest, at least when considering predator-prey interactions, are the second term in equation (1) and the first term in equation (2), called the *functional response* and *numerical response* respectively (Holling 1959). As we can see from the equation template, the functional response is a rate, specifically the rate of prey capture per predator. The simplest possible assumption is that the functional response is linear, or that the number of prey capture increases linearly with increasing numbers of prey. This simple assumption may be actually true over some ranges (Korpimäki and Norrdahl 1991) or when one is considering filter-feeding organisms, but more often than not is simply an

approximation. Increasing numbers of prey can create additional ecological interactions, not to mention changing the foraging behaviors of the predators. More realistic assumptions about the functional response have the rate of capture per predator decreasing with increasing number of predators. Even under this assumption, there are several different possibilities. For example, when prey are very abundant, predators will eventually become satiated. Another possibility is even more realistic: With very low numbers of prey, predators will lack the experience to be efficient hunters. With increasing numbers of prey, predators will become more efficient at hunting. Ultimately, there will be a number of prey beyond which the predators simply become satiated (Tinbergen 1960; Papaj and Lewis 1993).

Finally, the numerical response term correlates predator births to the number of prey captured. Because of this, the numerical response is itself a function of the functional response. Specifically, the numerical response depends on how many prey are in the population, how good the predators are at capturing them, and how much energy from the prey captures can be allocated to the production of new offspring. Naturally, this is a very complex term and will depend on other environmental variables, other stresses on the predator population, the energetic cost of offspring, etc. Ecologists almost always collapse most of this complexity in to a single parameter and represent the numerical response as a constant multiplied by the functional response.

Now that we have considered how the basic template could in principle be filled in, let's return to the Lotka-Volterra (L-V) model itself and to Volterra's justification of it. As we said earlier, the L-V model is probably the simplest way to make a population level predator-prey model because we are going to fill in (1) and (2) with the simplest functions.

In our representation of the L-V model, r stands for the growth rate of the prey population and m stands for the death rate of the predators. The functional response is linear, expressed as a constant a multiplied by V . Similarly, the numerical response is a linear function of the numerical response so the whole numerical response expression can be written as a parameter b multiplied by the functional response, or $b(aV)$. The L-V model is thus described with the following differential equations:

$$\frac{dV}{dt} = rV - (aV)P, \quad (3)$$

$$\frac{dP}{dt} = b(aV)P - mP. \quad (4)$$

These equations describe a model which predicts one result: the predator

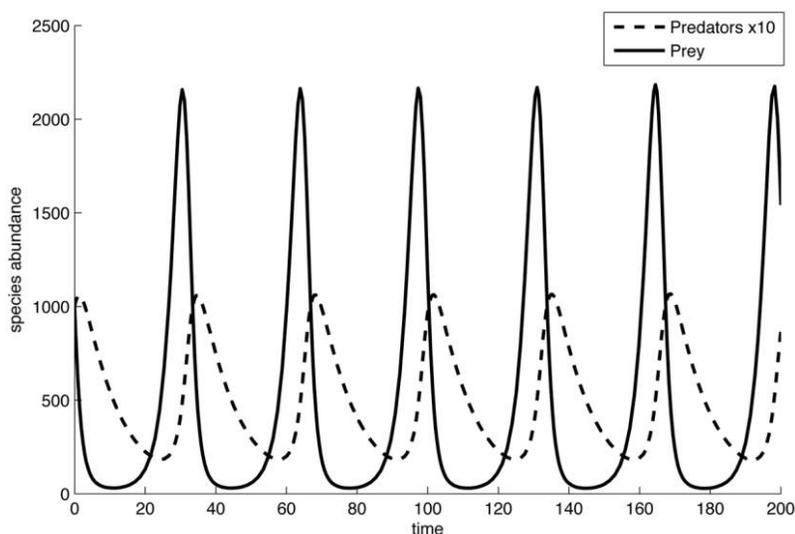


Figure 1. Lotka-Volterra model's oscillation.

and prey populations will oscillate indefinitely, out of phase with one another. Although for every set of parameter values with species coexistence, there exists one equilibrium where the populations do not oscillate, this equilibrium is unstable and hence a model population continues to oscillate if it is perturbed even slightly off of these equilibrium values.

Figure 1 plots the result of this oscillation for a set of parameter values and initial conditions. Qualitatively, it can be described as follows: As the prey population increases, the predator population increases as well, lagging behind. However, eventually, the predators begin to overtake the prey by continual feeding, which eventually begins to drive the prey population down in size. This results, in turn, in the predator population being driven down in size, and then the cycle repeats again. This undamped oscillation is the first important property of the L-V model that we will test for robustness in this paper. Before doing so, let us consider several more properties.

A second important property of the L-V model is neutral stability. The model exhibits neutrally stable oscillations, which means that perturbations away from the current oscillation amplitude will result in a new oscillation beginning from the point to which the oscillation was perturbed. There is no restoring force to bring the population back to the initial amplitude of the oscillations. This can be seen by plotting several trajectories, corresponding to different initial conditions, in the phase

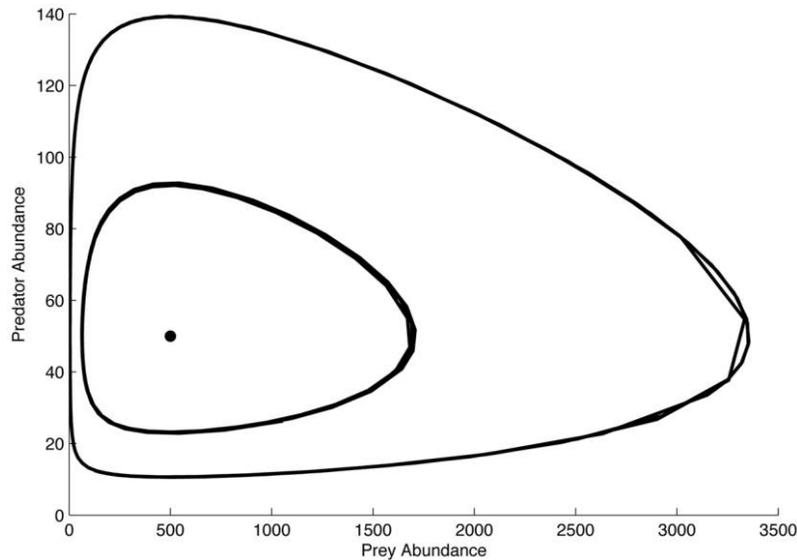


Figure 2. Phase-plot of Lotka-Volterra model.

space of the model (see Figure 2). Each loop is closed, corresponding to a stable oscillation. Perturbation results in the formation of a new loop, corresponding to a new, neutrally stable oscillation.

The third and most important property of the L-V model is what we will call the *Volterra Property*, which is the key component of the Volterra Principle. The Volterra Property states that a general biocide, any substance which has a harmful effect on both predators and prey, will *increase* the relative abundance of the prey population. To see this, we first need to solve for the equilibrium abundances of the species by setting each differential equation to zero. After some algebra, we find that the equilibrium values are

$$\hat{V} = \frac{m}{ab}, \quad (5)$$

$$\hat{P} = \frac{r}{a}. \quad (6)$$

These are unstable equilibria; however, they correspond to the average abundance of the predator and prey species over indefinitely long time periods.

We can derive the Volterra Property from the L-V model by first expressing the ratio of the average size of the predator population to the average size of the prey population (\hat{P}/\hat{V}) as ρ . Decreases in ρ will correspond to increases in the relative size of the predator population.

From equations (6) and (5) we can see that

$$\rho = \frac{rb}{m}. \quad (7)$$

The next step is to consider how a general biocide affects the model populations. We can represent the introduction of a biocide as corresponding to changes in r and m . Specifically, biocides decreases the prey growth rate (r) and increases the predator death rate (m). Inspecting ρ , the expression for the ratio of average densities, we can see that $\rho(\text{biocide}) < \rho(\text{normal})$ (Roughgarden 1997; May 2001, 439). Since smaller values for ρ mean a larger relative size of the prey population, the population of prey will increase relative to the number of predators when a biocide is applied. This is the Volterra Property: the general biocide increases the relative size of the prey population.

The Volterra Property is a key component of the Volterra Principle, the ecological theorem alleged to be robust. The second component of the Volterra Principle concerns a particular kind of causal structure: negative coupling. A predator species and a prey species are *negatively coupled* just in case increasing the abundance of predators decreases the abundance of prey and increasing the abundance of prey increases the abundance of predators.

The full formulation of the Volterra Principle connects this core causal structure to the robust property as follows:

Ceteris paribus, if a two-species, predator-prey system is negatively coupled, then a general biocide will increase the abundance of the prey and decrease the abundance of predators.³

The Volterra Principle has great ecological significance. One very practical consequence of it is that pesticides will often increase unwanted pests, such as the effect DDT had on citrus groves in southern California in the 1950s. When it was used as a treatment against scale insects, orchardists learned that DDT makes the pest problem worse. The cottony cushion

3. We know of no canonical formulation of the principle and there is considerable variation in the ecological literature. Our formulation is more general than the one in Weisberg 2006, which restricted the principle to systems where the abundance of predators is controlled mostly by the growth of the prey and the abundance of the prey controlled mostly by the death of predators. We believe that the principle holds more generally and only requires negative coupling between the two species.

scale insect (*Icerya purchasi*) population increased upon the application of DDT because along with the scale insects, the DDT killed the vedalia beetle (*Rodolia cardinalis*), a predator species keeping the scale insect under biological control (Elton 1958; Catagirone and Douth 1989).

The principle also has a deeper theoretical significance. It provides a vivid example of how ecologically coupled systems can behave in unexpected, nonlinear ways. Interventions in such systems are difficult and, without due care, can have the opposite effect of what was intended.

3. Parameter and Structural Robustness. The three key properties of the Lotka-Volterra model—undamped oscillations (1), neutral stability (2), and the Volterra Property (3)—are very interesting and would be of great ecological relevance if they were generally true of real populations. One way to ask whether these properties hold generally would be direct empirical investigation. We could go out in to the field or design a laboratory experiment and see whether these predictions are close to what really happens. But there is another approach to answering the question, one that is often conducted prior to or in conjunction with direct empirical investigation. This approach is robustness analysis. If we investigate related but distinct models, do we continue to see these three properties?

The first step in answering this question is to examine a representative sample of the models described by equations (3) and (4). One does this by evaluating the behavior of the model under different parameter settings. This analysis ensures that there is no special dependence of an interesting result on some particular set of initial conditions. We call this kind of analysis *parameter robustness analysis*, because it shows us whether the model's behavior is dependent on any particular set of parameters. Parameter robustness analysis is also known as *sensitivity analysis* in the modeling literature.⁴ Conducting parameter robustness analysis generates equivalence classes of parameter values that yield the same, or appropriately similar results.

Although we will not detail the analysis here, the three key properties of the L-V model are stable under wide ranges of parameter values for the L-V equations. In fact, as far as we know from surveying the literature and our own analyses, these three properties are robust across all parameter values where the two species coexist.

Parameter robustness can be applied alone when a theorist has good

4. There is an extensive literature on sensitivity analysis in population biology and modeling more generally. The basics are discussed in Grimm and Railsback 2005, Chapter 9. More extensive discussions of biological applications of sensitivity analysis can be found in Bartell et al. 1986 and Dreschler 1998. For more general discussions, see Rose 1989, Vose 2000, and Saltelli et al. 2004.

reason to believe that the basic structure of her model is adequate for her purposes, but remains uncertain about the values that should be assigned to the parameters. A second kind of robustness analysis relevant to our discussion of the Volterra Principle is what we call ‘structural robustness analysis’. In this kind of robustness analysis, the theorist considers changes to the causal structure of the system being modeled by analyzing models with a different mathematical structure. For example, in our model we might start with the L-V model, but add in terms representing predator satiation, the ability of prey to seek cover, multiple sources of food for the predator, or even complex adaptive behaviors such as learning. Any ecological interaction could in principle be added to the model.

Structural robustness analysis allows the theorist to address a different set of uncertainties from the ones addressed with parameter robustness. Using this approach, the theorist can probe which parts of the causal structure represented by her model are really essential for the production of an observed behavior of the model.⁵ This can take the form of adding new components to the causal structure, where one starts from a very minimal model such as the L-V model and adds new causal interactions. It can also involve removing factors, such as when one starts with a complex model, calibrated to a particular system, and removes factors to see which ones really make a difference. Structural robustness thus helps theorists isolate which components are necessary for the production of an important property, and which ones produce less robust properties.

In the next section, we describe one kind of structural robustness analysis of predation models. We consider a density-dependent version of the L-V model, which builds in an environmental resource that limits population growth. This model makes a major structural change to the L-V model, but keeps the core negative coupling intact. We focus on this robustness test because density dependence has played a particularly important role in the development of predation theory.

4. Predator-Prey Model with Density Dependence. While considerations of structural robustness could lead us to add any additional function to the predator-prey equations, a natural ecological addition would be to add a carrying capacity to the growth rate of the prey. If the predators did not exist, this carrying capacity indicates the maximum size to which the prey can grow, typically limited by resources in the environment. Adding a form of carrying capacity can be accomplished by making the prey population growth density dependent. A logistic growth term of the form $dV/dt = r(1 - V/K)$ is substituted for the first term in the prey equa-

5. This aspect of robustness analysis is discussed especially clearly by Odenbaugh 2007.

tion (Leslie 1948; Roughgarden 1997; Berryman 1992), yielding the following equations:

$$\frac{dV}{dt} = r \left(1 - \frac{V}{K} \right) V - (aV)P, \quad (8)$$

$$\frac{dP}{dt} = b(aV)P - mP. \quad (9)$$

In the model described by these equations, there are three equilibria, which correspond to the three possible outcomes in the long run. The first equilibrium is extinction of both species. The second equilibrium involves predator extinction, but the prey continues to survive and grow to its carrying capacity. The third equilibrium is of most interest to us and says that both species can coexist. Solving these equations for this third, coexistence equilibrium, we get the following expressions:

$$\hat{V} = \frac{m}{ab}, \quad (10)$$

$$\hat{P} = \frac{r}{a} \left(1 - \frac{m}{abK} \right). \quad (11)$$

Several things are worth noting about this equilibrium point. First and most important, it is a stable equilibrium; there are trajectories leading from the points in the vicinity of this equilibrium to this equilibrium. Once the populations settle on this point, they will not fluctuate in size unless they are perturbed. Population sizes in the vicinity of the other equilibrium values (full extinction or predator extinction) will likewise settle down to their respective equilibrium values.

The stability of this equilibrium can be demonstrated both analytically and graphically. Analytically, we can see that the equilibrium is stable by computing the eigenvalues of the Jacobian matrix. For all parameter sets that bring the population to the third equilibrium, the eigenvalues have a real part and the real part is negative. This corresponds to a stable equilibrium point (Roughgarden 1997; May 2001).

A graphical analysis is shown in Figures 3 and 4, which are graphs of the phase space for the density-dependent model with different values of K . In Figure 3, all of the trajectories can be seen leading in to the equilibrium point. In Figure 4, the trajectories spiral in to the point. The difference between these graphs is controlled by the magnitude of K . For smaller values of K , there is no oscillatory tendency in the approach to

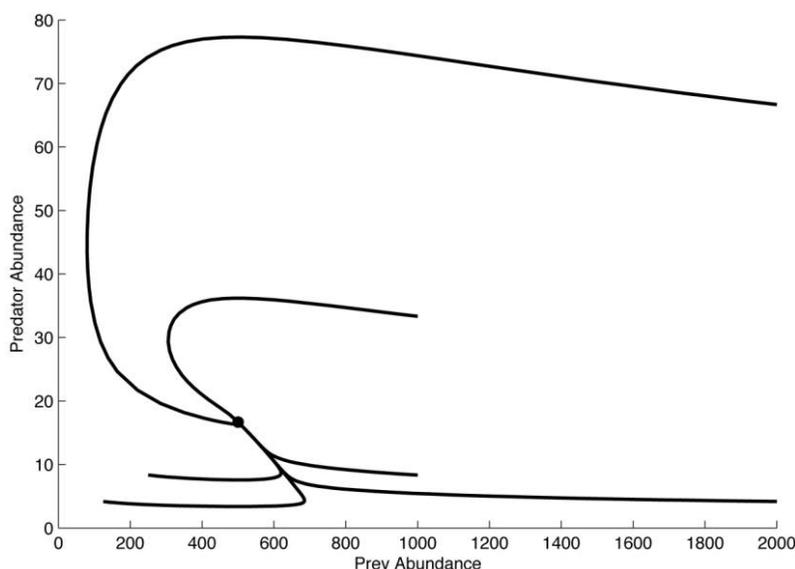


Figure 3. Density-dependent Lotka-Volterra model with $K = 750$.

the equilibrium point. When K is larger, such as in Figure 4, the trajectories have an oscillatory tendency as they approach the equilibrium point. The oscillatory tendency increased in magnitude as K is increased. As $K \rightarrow \infty$, the model becomes the L-V model and, as we would expect, the oscillations become undamped.

We are now in the position to make two comparisons between the density dependent model and the L-V model. The first property of the L-V model was undamped oscillations. In regions of state space where the density dependent model predicts coexistence, the model shows either no oscillations at all or damped oscillations, all leading to a stable equilibrium. Even an arbitrarily small amount of density dependence will destroy the undamped oscillation. Thus the first property of the L-V model is not structurally robust and cannot be formulated as part of a robust theorem. Since the first property is not robust, the second property—neutral stability in the oscillations—cannot be robust either.

To examine the third property of the L-V model, the Volterra Property, we once again express the coexistence equilibrium values as the ratio ρ , yielding

$$\rho = \frac{\hat{P}}{\hat{V}} = \frac{r(abK - m)}{aKm}. \quad (12)$$

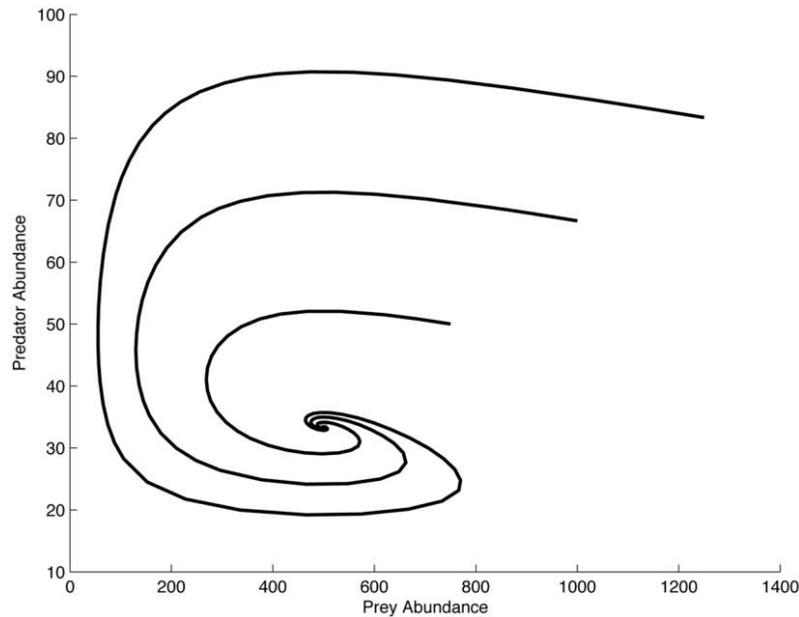


Figure 4. Density-dependent Lotka-Volterra model with $K = 1,500$.

Since r is in the numerator and m is in the denominator, we can see that the Volterra Property, and by extension the Volterra Principle, holds for this model. If we increase the death rate of the predators and decrease the growth rate of the prey, this corresponds to smaller values of ρ , meaning the relative size of the prey population increases. In fact, in this model the Volterra Principle has an even more direct interpretation. In the L-V model, the equilibrium values corresponded to the average size of each population. However in the density-dependent model, the terms in ρ are the actual equilibrium abundances of the two populations. Decreasing ρ will have a direct effect on the equilibrium size of the populations, not the average size over time.

Further structural robustness analysis would consider other changes to the causal structure represented in the model drawn from the kinds of ecological factors known to be relevant to population dynamics and predation. While any change to the basic structure is a kind of structural robustness test, ecologists are most interested in the ones that are potentially ecologically realizable. When a robust property survives all or some range of structural robustness tests, then we can say that the property is structurally robust to such and such changes to the causal structure of the system. If these changes sample a sufficiently broad set of ecologically

plausible circumstances, then ecologists will often simply refer to a phenomenon as robust.

5. Representational Robustness. Testing the structural robustness of a theorem is a matter of iteratively varying the basic assumptions of a model to see whether the theorem continues to hold. We could make many other modifications to the density-dependent predator-prey model to test the structural robustness of the L-V model. For example, we might examine other equations for the functional and numerical responses, examine the effects of population stochasticity, include the possibility of predator satiation, and add terms describing the prey's ability to hide from the predators. These possibilities were explored and shown to be robust in the classical ecology literature about predation.⁶ Instead of following those analyses here, we now turn to a different kind of robustness, which varies what we will call the *representational framework* of the model.

Mathematical models can be thought of as being composed of *state variables*, which are variables that represent the properties (states) of interest to the modeler and *transition rules*, the rules that govern how the states change through time (Lewontin 1963). The representational framework of the model is a general description of the type of state variables and the type of transition rules the model employs. For example, the variables in a biological model might represent individuals or populations. They might also represent more abstract properties of target systems, such as energy or nutrients, or even pathways by which these properties flow. Transition rules can be deterministic, probabilistic, or stochastic. They can also be discrete or continuous with respect to time.

The models we have considered so far use population densities as their state variables and have deterministic transition rules that are continuous with respect to time. Thus they were formulated using a set of differential equations. Versions of these models that use difference equations could very easily be generated, which would make them discrete with respect to time but hold the other aspects of the representational framework constant. Further changes can be made by adding in probabilistic transition rules or changing the perspective of the state variables from populations to energy flows or to individual organisms.

Parameter and structural robustness analysis vary assumptions about the target phenomena to see how the emergence of the phenomenon is sensitive to these assumptions. Representational robustness analysis lets us probe a different feature of our models. It holds these assumptions

6. See Roughgarden 1979 for a review of some of the key classical models.

fixed, but analyzes whether or not the way these assumptions are represented make a difference to the production of a property of interest.

Recent ecological literature has been especially concerned with the representational robustness of classic population based models when they are reconstructed in an *individual-based* framework, where the state variables are attached to individual organisms (Grimm and Railsback 2005). In the next sections, we illustrate representational robustness analysis by considering two novel individual-based models of predation with which we tested the Volterra Principle for representational robustness.

6. Individual-Based Predator-Prey Models. The models considered so far have treated predation as a population-level phenomenon. These models aggregate the properties of many organisms and represent them using just a handful of population-level state-variables. They contain no explicit representation of individuals or their properties, only the statistical aggregates of those properties. In contrast, individual-based models (IBMs) explicitly represent individuals and their properties. An IBM includes a set of state variables for each individual within the model population. It also include assumptions about how individuals in the population behave, develop, and interact over time. Since IBMs often contain thousands of variables, their dynamic consequences are usually investigated via computational simulation rather than mathematical analysis.

Population-level models are often more elegant and amenable to mathematical analysis than IBMs, but their very simplicity can be limiting. Ecological systems have rich structure that is not readily visible from a population-level perspective. Organisms within a population generally differ in their properties and life histories. Interactions between organisms are *local*, involving a few individuals at a particular place and time. IBMs are effective at capturing this individual-level detail because they explicitly represent the properties of each organism in a population, and because individuals can be set up to interact in small numbers on a spatial lattice. When one wants to test whether individual variability or local interactions affect the robustness of a generalization, one can build an IBM.

IBMs are also useful because they integrate our understanding of the different levels in the biological hierarchy. Ecologists gather data about organisms, populations, and communities. IBMs help reconcile these multiple levels of data, because assumptions about organisms and their interactions enter into an IBM, and the population or community consequences of these assumptions result from running IBM simulations. IBMs are thus informed and constrained by ecological data at multiple levels. In contrast, the points of contact between population-level models and data are all at the population and community levels. Such models include assumptions about populations, but they are either silent or ambiguous

concerning assumptions about individuals. This can be a strength when we want to remain agnostic about those assumptions, but it can also be a liability. In many cases, we may want to vary those assumptions and understand their consequences.

In recent years, IBMs have become increasingly common within ecology and among the sciences more generally. (See also Donalson and Nisbet 1999; DeAngelis and Mooij 2005; Grimm and Railsback 2005.) They are not a substitute for population-based models, but they can be used to relax assumptions made by generating population-based models. As such, each of these frameworks may be more or less appropriate depending on the purpose at hand. Indeed, for the purpose of testing the robustness of a generalization, one ought to examine as many representational frameworks as possible.

To test the representational robustness of the three key properties of the L-V model, we will translate its variables, parameters, and other assumptions into individual-based terms.⁷ In our discussion of the L-V model, we showed how the model makes assumptions about the growth and death rates of the predator and prey populations and about the interaction between these populations in the form of predation. Any individual-based version of this model must reconceive these processes in terms of individuals; it must make assumptions about the births and deaths of discrete predator and prey individuals, and the interactions between these discrete individuals in the form of predation.

In our first IBM version of the L-V model, we assume that individuals move about on a 30×30 toroidal lattice composed of 900 cells. Each individual has three variables: a binary variable denoting whether the individual is predator or prey, and two integer variables denoting a vertical and horizontal position on the lattice. Time is discrete; a global clock advances one tick at a time. For each tick of the global clock, all individuals execute a fixed set of rules that determine how they move on the lattice, reproduce, die, and interact with others. The rules for predators are as follows:

Movement rule. Move one step in a random direction.

Predation rule. Check if there are any prey on the current cell. If so, select one at random, catch it, and pick a random number from 1

7. The models described using this section and the next were developed in NetLogo v. 3.0.2 (Wilensky 1999). The source code for the models is available at <http://philsci-archive.pitt.edu>. Our models have some similarities to another NetLogo model called 'Wolf Sheep Predation' (Wilensky 1998). For more information on 'Wolf Sheep Predation', see Wilensky and Reisman 2006.

to 100. If this number is less than or equal to the parameter *predator-conversion* then reproduce.

Death rule. Pick a random number from 1 to 100. If this number is less than or equal to the parameter *predator-death-probability* then die.

These rules, when executed by each predator on the lattice, correspond roughly to assumptions made in the L-V model, but notice that these rules are not determined by that model. To translate any population-based model into individual-based terms, we must make explicit assumptions about individuals that were either implicit or undefined in the population-based version. This means that there is typically no uniquely correct way to carry out the translation from population-based to individual-based models.

The IBM assumes that predators move randomly on a two dimensional toroidal lattice. The L-V model, on the other hand, makes no assumption about movement at all. It is consistent with the assumptions that all individuals move, that some individuals move, or even, strictly speaking, that no individuals move. It places no explicit constraints on what intrinsic or environmental factors determine movement or even whether the predator and prey move in a probabilistic or deterministic fashion.

Moreover, the IBM assumes that predators catch prey by randomly selecting one prey individual from all that are located on the same cell. Once again, this is one of the many possible assumptions we could have made to develop an IBM analogue of the L-V model. We could have represented predation without using a spatial lattice, where predators randomly choose prey individuals from the whole prey population. We could also have used a different predation rule on a lattice. For example, the predation rule could have stated “if a predator is within 1 cell of a prey, then the prey is consumed.” The L-V model does not strictly correspond to any of these particular assumptions. A modeler who wishes to construct an IBM, however, must make an explicit decision about them.

Our corresponding rules for the prey are as follows:

Movement rule. Move one step in a random direction.

Reproduction rule. Pick a random number from 1 to 100. If this number is less than or equal to the parameter *prey-reproduction-probability* then reproduce.

Death rule. Check if I have been caught by a predator. If so, then die.

Together, the predator and prey rule-sets comprise one possible IBM interpretation of the L-V model. Notice that, as with the L-V model, this

IBM defines a negatively coupled predator-prey system. *Ceteris paribus*, increasing the abundance of predators in the IBM will increase the chance that prey are captured, and thus decrease the expected number of prey. *Ceteris paribus*, increasing the abundance of prey will increase the chance of a predation instance. Since each such instance has a fixed probability of resulting in predator reproduction, increasing the abundance of prey will increase the expected number of predators. It is critical that the IBM does indeed define a negatively coupled system, because negative coupling is a necessary condition for a system to demonstrate the Volterra Principle.

To understand the dynamic consequences of this model, we set up a computational representation of a lattice, place predator and prey individuals on the lattice, iteratively execute the predator or prey rules for each individual, and observe how the system evolves over time. In the initial state used in our simulations, there are V prey and P predators positioned randomly on the lattice.

After observing many simulations with differing parameter sets and initial conditions, we concluded that this IBM does not exhibit stable oscillations in the numbers of predators and prey, the first property of the L-V model.⁸ There are sets of parameters which initially result in oscillations, but these oscillations are unstable, increasing in amplitude over time until either both species have gone extinct (Figure 5), or else the predators have gone extinct and only the prey remain (Figure 6). Because this is a probabilistic model, the same parameter set and initial conditions sometimes results in two-species extinction and sometime results in the prey surviving. Figures 5 and 6 correspond to this situation: both used the same parameter set and initial conditions.

Regardless of the parameters, one or both species invariably goes extinct. This result suggests that stable oscillations, or even stable coexistence, is not a representationally robust feature of predator-prey systems. Since the IBM does not exhibit stable oscillations, it clearly will not exhibit property 2 of the L-V model, neutrally stable oscillations.

Finally, since the IBM does not exhibit coexistence of species, even in the short or medium term, it cannot exhibit the Volterra Property. Testing this property requires examining the effect of a biocide on average abundances, which are either $\hat{P} = 0$ and $\hat{V} = \infty$ or $\hat{P} = 0$ and $\hat{V} = 0$ for this model. Coexistence of predator species is a precondition for the Volterra

8. The model was analyzed in the following ranges of parameter values and initial conditions: All runs had *pred-conversion-prob* = 0.5, *prey-conversion-prob* = 0.2 and ran until the predators died out (always fewer than 3,000 cycles). We varied the following parameters *prey-reproduction-prob* varied between 0.05 and 0.15, *pred-death-prob* varied between 0.05 and 0.09, *pred-initial-number* varied between 50 and 500, and *prey-initial-number* varied between 50 and 500.

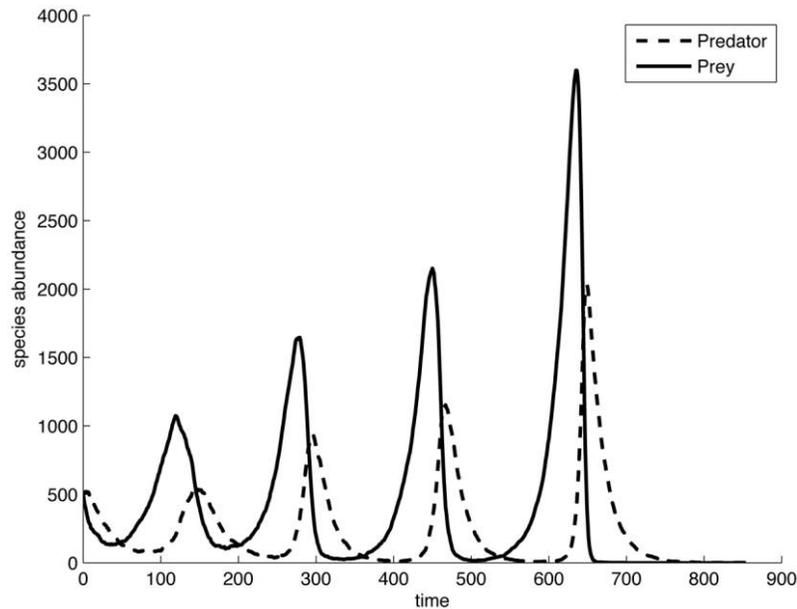


Figure 5. Simple individual-based model corresponding to the Lotka-Volterra model. In this trial, both species go extinct.

Principle to hold. To test the representational robustness of this property and the Volterra Principle itself, we must begin with a predator-prey model which has quasistable behavior for a reasonable length of time.

Since we would like to test for the robustness of the Volterra Principle, we need to find an IBM that exhibits coexistence of the two species. In the next section, we describe a modification that does stabilize the populations. It achieves this stabilization by adding density-dependence.

7. Density Dependent IBM. One source of instability in the previous model is the lack of limits on population growth. The population oscillations tend to become more pronounced with each successive cycle until the population of one or both of the species falls too low and the species goes extinct. This suggests that if we introduce a carrying capacity to limit the upward amplitude of the oscillations in this model, extinction will be less likely.

How can we impose a carrying capacity on the model? The most direct way is to assume a fixed global limit on the number of predators and prey, yet this is not in the spirit of individual-based modeling; it is a top-down assumption about populations, rather than a bottom-up assumption

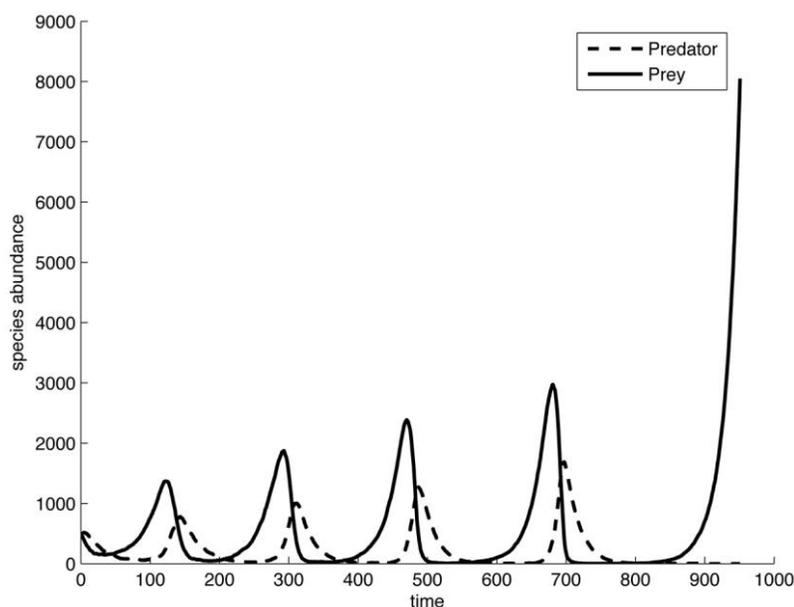


Figure 6. Simple individual-based model corresponding to the Lotka-Volterra model. In this trial, only the predators go extinct.

about individuals. A bottom-up alternative is to impose a limited resource into the model, such as space. For example, we might assume that at most one predator can occupy any cell in the lattice at a given time.⁹ Another type of limited resource is food for the prey population. For example, if the prey are herbivores, the limited availability of edible foliage in the environment imposes a natural carrying capacity on the prey.

In our density-dependent IBM, we assume that the size of the prey population is limited by availability of food in the environment (for convenience, we will call the food ‘foliage’, but it could represent any naturally available resource). We assume that each cell of the lattice either contains a unit of foliage or not. When eaten by a prey individual, the unit of foliage disappears, and it then has a certain probability (set by the parameter *foliage-growth-prob*) of reappearing at any subsequent tick. These assumptions suggest a revised rule-set for prey with a new foliage rule and a revised death rule:

9. A limitation to this assumption is that carrying capacity would be directly linked to the size of the lattice, so there would be no independent way to vary lattice size and carrying capacity.

Movement rule. Move one step in a random direction.

Foraging rule. Check if there is foliage on the current cell. If so, eat it, and pick a random number from 1 to 100. If this number is less than or equal to the parameter *prey-conversion-probability* then reproduce.

Death rule. Check if I have been caught by a predator. If so, then die.

The rule-set for predators remains the same.

Does this modified IBM, with density-dependence, display any of the three basic properties of the L-V model discussed earlier? After investigating many different initial states and parameter sets, we concluded that there is a wide range of parameter conditions for which this model does exhibit oscillations in the numbers of predators and prey for very long periods of time.¹⁰ The amplitude of the population oscillations tends to vary stochastically over time, but both species do persist and oscillate for long runs of the model.¹¹ Thus, the density-dependent IBM exhibits an analogue of property 1 of the L-V model; for very long time intervals, it exhibits undamped oscillations.¹²

The model does not appear to exhibit property 2, neutrally stable oscillations. Under all the conditions we examined which result in stable oscillations, the average abundance of predators and prey did not depend upon initial conditions or prior population sizes. After perturbing the populations away from their equilibrium temporal average sizes, the populations would always return to their former averages. This suggests that property 2 of the L-V model is not representationally robust.

To check whether the model exhibits the Volterra Property, we must somehow simulate the effect of a general biocide that would elevate the death rate of both the predators and of the prey. Since there are no parameters in the model that correspond directly to these rates, we must manipulate them indirectly. Fortunately, the individual-based framework makes it easy to simulate the dispersion of a general biocide into our model system. We performed the following perturbation: First, initiate a typical simulation of the predator-prey system and wait long enough for

10. The same parameter sets were investigated for this model as in the non-DD IBM. We set the *foliage-growth-probability* to 0.05 for these simulations.

11. Since this is a stochastic model, both species will go extinct in the long run with probability one. However, we examined the model for very long runs (1×10^6 ticks) and observed oscillations. If the frequency of the oscillations is calibrated to the famous Lynx-Hare predator prey system, then this is equivalent to about 100,000 years.

12. Technically, these oscillations are called *long-lived transient oscillations*.

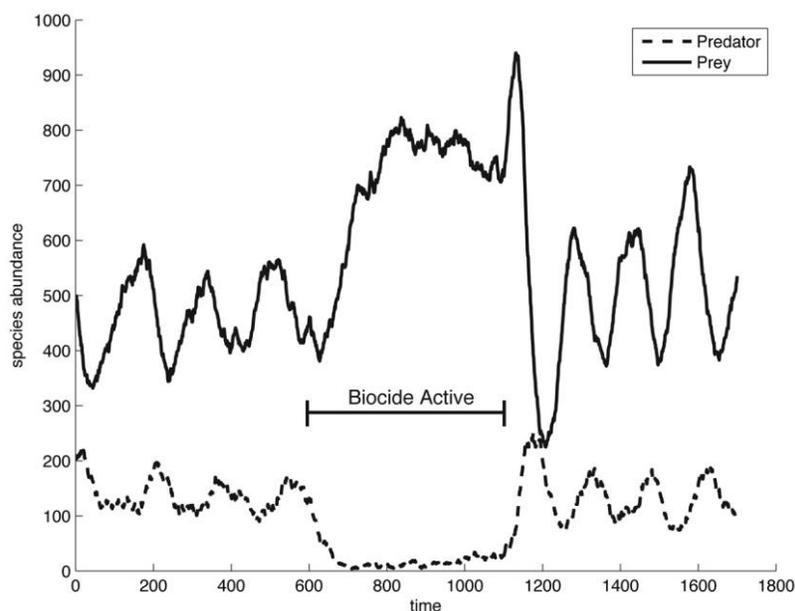


Figure 7. Density-dependent individual-based predation model with biocide perturbation.

the temporal average size of each population to reach a steady state. Next, randomly select some cells on the lattice to become ‘poisonous’, so that any predator or prey that lands on the cell will die. Since movement is random, predators and prey are equally likely to die as a result of landing on poisonous cells and the result is an increase in the death rate of both populations. Finally, wait for the temporal average size of each population to reach a new equilibrium.

After performing this perturbation over a broad range of parameter settings,¹³ we found that introduction of a general biocide tended to increase the average size of the prey population and to decrease the average size of the predator population (Figure 7). Thus, this revised IBM does exhibit the Volterra Property. Moreover, since this model also defines a negatively coupled predator-prey system, it satisfies the Volterra Principle. In other words, despite moving from a population to an individual-based framework, and despite altering various assumptions of the L-V model,

13. Again, we used the same set of parameter settings, but ran the simulation for 3,000 cycles, with the ‘poison’ interval between cycles 1,001 and 2,000. We set *biocide-abundance* to 0.018 for these simulations.

the Volterra Principle still held up. This suggests that it is representationally robust.

The analyses in this paper show that the Volterra Principle exhibits three kinds of robustness. Naturally, the scope of the Volterra Principle's robustness is not a settled matter. Analysis could continue with different and more realistic models. For example, we might investigate a model where reproduction and death are tied only indirectly to food consumption via the introduction of rudimentary elements of metabolism.¹⁴ In fact, tests for robustness are always an ongoing matter. It is unlikely that a theoretical community will settle the issue of a theorem's robustness once and for all.

8. Conclusion. Volterra discovered the principle which bears his name in 1926. By constructing a series of models, all similar but differing in some respects, ecologists have shown that the Volterra Principle is robust and accurately describes a real ecological phenomenon—when two species in a negatively coupled, predator-prey relationship coexist, a general biocide will favor the prey over the predators. In our discussion of four models of predation, we have outlined three kinds of robustness analysis which correspond to three different levels at which a result's robustness can be determined in modeling. These are parameter robustness, a result's surviving changes to the parameter set of a dynamical model; structural robustness, a result's surviving changes to the mathematical structure of the model; and finally, representational robustness, a result's surviving changes to the whole representational framework in which the model has been framed. Like all of the most important robust theorems, the Volterra Principle possesses all three kinds of robustness.

Robustness analyses of the Volterra Principle carried out by the ecological community, as well as the novel representational robustness analysis carried out in this paper, has several implications. Most importantly, robustness analysis has shown that the principle is highly general and will hold under a wide variety of conditions. It is not dependent on idealizing assumptions made in various models of predation. While any given model contains idealizing assumptions, analysis across models has allowed us to control for them and factor them out. The principle is also insensitive to many other detailed assumptions made in the modeling process (for example, concerning the mode of animal movement, reproduction, metabolism, etc.) This insensitivity to detail helps explain why the Volterra Principle has been confirmed in widely disparate natural systems, from

14. We developed such a model and tested it successfully for the Volterra Principle. As with the other individual-based models, it can be found at <http://philsci-archive.pitt.edu>.

cottony cushion scale insects and vedalia beetles to cod, Norwegian lobster, and sharks. In contrast, robustness analysis has revealed that several other properties of the original Lotka-Volterra model, including stable oscillations and neutral stability, are not robust in this respect.

Levins (1966, 20) wrote that “our truth is at the intersection of independent lies.” We do not see models, even the highly idealized ones involved in the robustness analysis of the Volterra Principle as ‘lies’, yet we think Levins’ point is correct. When studying phenomena as complex and hard to measure as predation, scientists often have little choice but to build approximate, idealized models. Finding that some result is robust across these models, however, is an important step in the process of a theorem’s confirmation.

REFERENCES

- Bartell, S. M., J. M. Breck, R. H. Gardner, and A. L. Brenkert (1986), “Individual Parameter Perturbation and Error Analysis of Fish Bioenergetics Models”, *Canadian Journal of Fisheries and Aquatic Sciences* 5: 160–168.
- Berryman, A. A. (1992), “The Origins and Evolution of Predator-Prey Theory”, *Ecology* 73: 1530–1535.
- Briggs, C. J., and M. F. Hoopes (2004), “Stabilizing Effects in Spatial Parasitoid-Host and Predator-Prey Models: A Review”, *Theoretical Population Biology* 65: 299–315.
- Catagirone, L. E., and R. L. Doutt (1989), “The History of the Vedalia Beetle Importation to California and Its Impact on the Development of Biological Control”, *Annual Review of Entomology* 34: 1–16.
- DeAngelis, D. L., and W. M. Mooij (2005), “Individual-Based Modeling of Ecological and Evolutionary Processes”, *Annual Review of Ecology, Evolution, and Systematics* 36: 147–168.
- Donalson, D. D., and R. M. Nisbet (1999), “Population Dynamics and Spatial Scale: Effects of System Size on Population Persistence”, *Ecology* 80: 2492–2507.
- Dreschler, M. (1998), “Sensitivity Analysis of Complex Models”, *Biological Conservation* 86: 401–412.
- Elton, C. S. (1958), *The Ecology of Invasions by Animals and Plants*. New York: Wiley.
- Forber, P. (2007), “On Biological Possibility and Confirmation”, manuscript. Medford, MA: Tufts University.
- Grimm, V., and S. F. Railsback (2005), *Individual-Based Modeling and Ecology*. Princeton, NJ: Princeton University Press.
- Hanski, I., H. Henttonen, E. Korpimäki, L. Oksanen, and P. Turchin (2001), “Small-Rodent Dynamics and Predation”, *Ecology* 82: 1505–1520.
- Holling, C. S. (1959), “The Components of Predation as Revealed by a Study of Small Mammal Predation of the European Pine Sawfly”, *Canadian Journal of Entomology* 91: 293–320.
- Korpimäki, E., and K. Norrdahl (1991), “Numerical and Functional Responses of Kestrels, Short-Eared Owls, and Long-Eared Owls to Vole Densities”, *Ecology* 72: 814–826.
- Leslie, P. H. (1948), “Some Further Notes on the Use of Matrices in Population Analysis”, *Biometrika* 35: 213–245.
- Levins, R. (1966), “The Strategy of Model Building in Population Biology”, in E. Sober (ed.), *Conceptual Issues in Evolutionary Biology*. Cambridge, MA: MIT Press, 18–27.
- Lewontin, R. C. (1963), “Models, Mathematics, and Metaphors”, *Synthese* 15: 274–296.
- Lotka, A. J. (1956), *Elements of Mathematical Biology*. New York: Dover.
- MacArthur, R. H., and J. H. Connell (1966), *The Biology of Populations*. New York: Wiley.

- May, R. M. (2001), *Stability and Complexity in Model Ecosystems*. Princeton Landmarks in Biology. Princeton, NJ: Princeton University Press.
- Maynard Smith, J. (1974), *Models in Ecology*. Cambridge: Cambridge University Press.
- Odenbaugh, J. (2007), "True Lies: Robustness and Idealizations in Ecological Explanations", manuscript. Portland, OR: Lewis and Clark College.
- Orzack, S. H., and E. Sober (1993), "A Critical Assessment of Levins's *The Strategy of Model Building in Population Biology* (1966)", *Quarterly Review of Biology* 68: 533–546.
- Papaj, D., and A. Lewis (1993), *Insect Learning*. New York: Chapman & Hall.
- Ricklefs, R. E., and G. L. Miller (2000), *Ecology*. New York: Freeman.
- Rose, K. A. (1989), "Sensitivity Analysis in Ecological Simulation Models", in M. G. Singh (ed.), *Systems and Control Encyclopedia*. New York: Pergamon, 4230–4235.
- Roughgarden, J. (1979), *Theory of Population Genetics and Evolutionary Ecology: An Introduction*. New York: Macmillan.
- (1997), *Primer of Ecological Theory*. Upper Saddle River, NJ: Prentice-Hall.
- Royama, T. (1971), "A Comparative Study of Models for Predation and Parasitism", *Researches on Population Ecology*, Supplement 1: 1–91.
- Saltelli, A. F., F. Tarantola, F. Campolongo, and M. Ratto (2004), *Sensitivity Analysis in Practice: A Guide to Assessing Scientific Models*. New York: Wiley.
- Tinbergen, L. (1960), "The Natural Control of Insects in Pinewoods: I. Factors Influencing the Intensity of Predation by Songbirds", *Archives Néerlandaises de Zoologie* 13: 266–336.
- Volterra, V. (1926a), "Fluctuations in the Abundance of a Species Considered Mathematically", *Nature* 118: 558–560.
- (1926b), "Variazioni e fluttuazioni del numero d'individui in specie animali conviventi", *Memorie della R. Accademia Nazionale dei Lincei* 2: 5–112.
- Vose, D. (2000), *Risk Analysis: A Quantitative Guide*. Chichester: Wiley.
- Weisberg, M. (2006), "Robustness Analysis", *Philosophy of Science* 73: 730–742.
- Wilensky, U. (1998), "Netlogo Wolf Sheep Predation Model". Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University, <http://ccl.northwestern.edu/netlogo/models/WolfSheepPredation>.
- (1999), *Netlogo*. Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University.
- Wilensky, U., and K. Reisman (2006), "Thinking Like a Wolf, a Sheep or a Firefly: Learning Biology through Constructing and Testing Computational Theories—an Embodied Modeling Approach", *Cognition and Instruction* 24: 171–209.
- Wimsatt, W. C. (1981), "Robustness, Reliability, and Overdetermination", in M. Brewer and B. Collins (eds.), *Scientific Inquiry and the Social Sciences*. San Francisco: Jossey-Bass, 124–163.