

# Robustness Analysis

Michael Weisberg<sup>†‡</sup>

---

Modelers often rely on robustness analysis, the search for predictions common to several independent models. Robustness analysis has been characterized and championed by Richard Levins and William Wimsatt, who see it as central to modern theoretical practice. The practice has also been severely criticized by Steven Orzack and Elliott Sober, who claim that it is a nonempirical form of confirmation, effective only under unusual circumstances. This paper addresses Orzack and Sober's criticisms by giving a new account of robustness analysis and showing how the practice can identify robust theorems. Once the structure of robust theorems is clearly articulated, it can be shown that such theorems have a degree of confirmation, despite the lack of direct empirical evidence for their truth.

---

**1. Introduction.** Physical scientists can often construct and analyze models that are derived from fully developed, well-confirmed background theories. A scientist in possession of such a theory can be very confident about her ability to construct realistic models and to make accurate predictions with these models. Newtonian mechanics, for example, is so well developed that mathematical models derived from it have guided satellites to other planets in our solar system.

Unfortunately, many phenomena, such as those studied by population biologists, have yet to be described by comprehensive theories, and it is unlikely that such theories will be developed in the near future. There is a large body of theoretical work in ecology, evolution, and systematics, but there are few overarching background theories that unify these domains. Thus in population biology and other sciences that deal with highly

<sup>†</sup>To contact the author, please write to: Department of Philosophy, University of Pennsylvania, 433 Logan Hall, Philadelphia, PA 19104-6304; e-mail: [weisberg@phil.upenn.edu](mailto:weisberg@phil.upenn.edu).

<sup>‡</sup>I am grateful for the extremely helpful comments of Patrick Forber, Peter Godfrey-Smith, Richard Levins, Ryan Muldoon, Deena Skolnick Weisberg, Brian Skyrms, Elliott Sober, C. Kenneth Waters, as well as audiences in the Biology Department at Penn, the Philosophy of Science Association, and the International Society for History, Philosophy, and Social Studies of Biology.

Philosophy of Science, 73 (December 2006) pp. 730–742. 0031-8248/2006/7305-0024\$10.00  
Copyright 2006 by the Philosophy of Science Association. All rights reserved.

complex systems, theorists generate multiple models for individual phenomena, all of which are idealized.

If this situation occurred when dealing with simple physical systems, one of the fundamental physical theories could be used to determine how much distortion was introduced by each idealization. Such theories have the resources to estimate the effect of various idealizations, providing guidance about what must be included when particular degrees of accuracy and precision are required. In the study of many complex systems, however, such a theory is unavailable. Theorists thus require an alternative method for determining which models make trustworthy predictions and which models can reliably be used in explanations. *Robustness analysis* provides such an alternative method.

According to Richard Levins, robustness analysis can show us “whether a result depends on the essentials of the model or on the details of the simplifying assumptions” (1966, 20). This is accomplished by studying a number of similar, but distinct, models of the same phenomenon:

[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 the intersection of independent lies. (20)

Despite its importance in scientific practice, robustness analysis has received surprisingly little attention in the literature about scientific method. Although Levins drew attention to the idea 40 years ago in his well-known paper “The Strategy of Model Building in Population Biology” (1966), until recently, only William Wimsatt, Steven Orzack, and Elliott Sober have discussed the topic in detail. In a sympathetic reconstruction, Wimsatt explains the core idea behind robustness analysis:

[A]ll the variants and uses of robustness have a common theme in the distinguishing of the real from the illusory; the reliable from the unreliable; the objective from the subjective; the object of focus from artifacts of perspective; and, in general, that which is regarded as ontologically and epistemologically trustworthy and valuable from that which is unreliable, ungeneralizable, worthless, and fleeting. (1981, 128)

From Wimsatt, we learn that robustness analysis’s aim is to separate the scientifically important parts and predictions of our models from the illusory ones that are accidents of representations—something physicists do using their comprehensive theories.

Not all philosophers of science have been sympathetic to Levins’ idea of robustness analysis. In a review of “The Strategy,” Orzack and Sober (1993) argue that robustness analysis appears to be a form of nonempirical

confirmation and thus has no place in scientific practice. In this paper, I endeavor to answer Orzack and Sober's criticism of the efficacy of robustness analysis. This requires developing a more detailed, albeit preliminary, account of robustness analysis and the notion of a 'robust theorem'. I will argue that robustness analysis is effective at identifying robust theorems, and while it is not itself a confirmation procedure, robust theorems are likely to be true.

**2. Orzack and Sober on Robustness.** Levins' original discussion of robustness analysis provides little general characterization of the practice. Although he discusses several examples, the only abstract characterization that Levins gives is something like the following: When confronted with multiple models of the same phenomenon, the theorist needs a way to separate reliable predictions from artifacts of the assumptions made by the different models. Thus the theorist examines multiple models, preferably ones that make qualitatively different assumptions about the phenomenon, and looks for a common prediction among these models. This common prediction is the "truth at the intersection of independent lies," or a *robust theorem*.

Since Levins provides no further account of robustness analysis, Orzack and Sober begin their critique by providing an interpretation of Levins' article. The first issue they discuss is whether robustness analysis is supposed to be a form of confirmation or merely a heuristic of some kind. They conclude that it is intended to be a novel, nonempirical form of confirmation:

When [Levins] further writes (p. 427) that a particular "non-robust" theorem "cannot be asserted as a biological fact" it becomes clear that Levins means that a statement's *robustness*, as distinct from its *observational confirmation*, can be evidence for its truth. (Orzack and Sober 1993, 538)

Their second step toward giving a more formal account of robustness analysis involves defining the term 'robust theorem'. On Orzack and Sober's reading, robust theorems are literally theorems, propositions that are the logical consequence of each member of a set of models. Schematically, robustness analysis consists in analyzing a set of models  $M_1, \dots, M_n$  and showing that, for every  $i$ ,  $M_i \models R$ . Robustness analysis is thus the search for this type of implication, which is supposed to generate confirmation.

Since robustness analysis is conducted by examining models, not data, Orzack and Sober argue that this procedure is nonempirical. Hence they believe that Levins is promoting a distinct form of confirmation, one that does not involve examining data. Rightly skeptical of nonempirical con-

firmation, Orzack and Sober present an argument that is intended to show the ineffectiveness of such a procedure.

The argument begins by enumerating possible circumstances under which a theorist might engage in robustness analysis. Considered abstractly, Orzack and Sober argue, there are three possible circumstances: The theorist either knows that all the models in the set are false, knows that at least one is true, or does not know whether any of them are true. Robustness analysis, they argue, is sure to yield a true theorem only in the case in which we know ahead of time that at least one model in the set is true. In the other two cases, knowing that  $R$  is entailed by each  $M_i$  provides no evidence about the truth of  $R$ .

One can further develop Orzack and Sober's argument with the following line of reasoning: If we knew that all the models in the set were true, we would not need robustness analysis in the first place. Robustness analysis is supposed to help us precisely in situations in which all the models are idealized (i.e., false), and we want to know what results drawn from our models are reasonable to believe. Since robustness analysis is ineffective in this case, it can never be a useful tool.

In his response to Orzack and Sober's analysis, Levins accuses them of providing an analysis that "is not relevant to the problem that was posed" (1993, 533). He also argues that Orzack and Sober's characterization of robustness as a rival to ordinary empirical confirmation is mistaken. Observation, he tells us, is not forsaken in robustness analysis:

Observation enters first in the choice of the core model and the selection of plausible variable parts, and later in the testing of the predictions that follow from the core model. (554)

Further elaborating his original position, Levins tells us that "the search for robust theorems reflects the strategy of determining how much we can get away with not knowing, and still understand the system" (554).

In these passages, Levins is characterizing robustness analysis quite differently than Orzack and Sober have. Levins does not believe that he is offering an alternative to empirical confirmation, yet he is vague about the exact nature of robustness analysis. He says that it concerns "how much we can get away with not knowing" while "still understand[ing]" biological phenomena, but does not tell us how it helps to confirm models and their predictions.

Despite sympathizing with Orzack and Sober's desire for a more detailed account and defense of robustness, I believe that their analysis is too abstract to be used for this purpose. While I do not dispute the validity of their argument for arbitrary sets of models and arbitrary logical consequences of those models, this result does not necessarily apply to specific subsets of models or to particular kinds of logical consequences of these

models. In fact, I will argue in Section 4 that robustness analysis can play a role in confirming hypotheses because of the structure of these hypotheses.

Despite this problem, Orzack and Sober's criticisms helpfully suggest a way to begin analyzing robustness analysis. Like them, I believe that the first step toward analyzing robustness is to clarify the concepts of 'robust theorem' and 'robustness analysis'. Only after we understand the nature of these concepts can we consider the rationality of employing robustness analysis.

**3. Discovering Robust Theorems.** The structure of robust theorems and the nature of robustness analysis are best appreciated through an examination of real scientific examples, not abstract philosophical analysis. Models of predator-prey interactions have been thoroughly studied by ecologists, and this study has yielded several robust as well as many nonrobust theorems. This makes predator-prey models excellent examples with which to study robustness analysis. In Section 3.1, I will describe how theorists have analyzed models of predator-prey systems to discover a robust theorem called the Volterra principle. This will form the basis of a more general characterization of robustness analysis that I give in Section 3.2.

*3.1. The Volterra Principle.* The simplest predator-prey model is known as the Volterra model or the Lotka-Volterra model, as it was independently analyzed by Alfred Lotka (1956) and Vito Volterra (1926a, 1926b). In this model,  $V$  represents the size of the prey population and  $P$  the size of the predator population. The parameter  $r$  stands for the intrinsic growth rate of the prey population,  $m$  is the intrinsic death rate of the predators, and the other parameters ( $a$  and  $b$ ) correspond to the prey capture rate and the rate at which each predator converts captured prey into more predator births. The model is described with the following differential equations:

$$\frac{dV}{dt} = rV - (aV)P, \quad (1)$$

$$\frac{dP}{dt} = b(aV)P - mP. \quad (2)$$

This model exhibits a regular, periodic oscillation in the size of both the predator and prey populations. Since there is no stable equilibrium described by these equations, the model predicts that the populations will oscillate indefinitely. In addition, the amplitude of the oscillations (peak

size of the populations) depends only on the initial conditions, not on the values for the parameters in the model. These two features are not robust, and a slight change to the model, such as the introduction of density dependence, will cause them to vanish. However, there is one important property of the model that is robust. This property can be seen when we consider how a general pesticide, an intervention that kills both predators and prey, affects the populations. Further analysis is required to discover this effect.

Although equations (1) and (2) describe a model with no stable equilibrium, there is one unstable equilibrium that corresponds to the time-averaged size of the two model populations (Hofbauer and Sigmund 1998, 16). We can find this equilibrium by setting each differential equation to zero and doing some algebra:

$$\hat{P} = \frac{r}{a}, \quad (3)$$

$$\hat{V} = \frac{m}{ab}. \quad (4)$$

We now need to examine what will happen when both populations are affected by a general pesticide. The easiest way to do this is to consider what happens to the ratio ( $\rho$ ) of the average size of the predator population to the average size of the prey population.

From equations (3) and (4), we can see that

$$\rho = \frac{\hat{P}}{\hat{V}} = \frac{rb}{m}. \quad (5)$$

The next step is to consider how the general pesticide affects this ratio. If we treat the pesticide as roughly equivalent to decreasing the intrinsic growth rate of the prey and increasing the intrinsic death rate of the predators, we find that  $\rho(\text{pesticide}) < \rho(\text{normal})$ . Since smaller values for  $\rho$  mean a larger relative size of the prey population, the population of prey will increase relative to the number of predators upon application of the pesticide. This result is known as the *Volterra principle* (Roughgarden 1979, 439). In order to see whether this property is robust, we will need to examine other models of predation.

Since some of the other properties of the Volterra model are destroyed by the introduction of density dependence, one obvious model to consider adds density dependence to the simple Volterra model. This can be ac-

completed by adding to equation (1) a term for the prey population's carrying capacity. The model can then be represented as follows:

$$\frac{dV}{dt} = r \left( 1 - \frac{V}{K} \right) V - (aV)P, \quad (6)$$

$$\frac{dP}{dt} = b(aV)P - mP. \quad (7)$$

After solving once again for the equilibrium values, which gives us the average population sizes, we can express the ratio  $\rho$  as

$$\rho = \frac{\hat{P}}{\hat{V}} = \frac{r(abK - m)}{aKm}. \quad (8)$$

By inspection, we can see that, once again,  $\rho(\text{pesticide}) < \rho(\text{normal})$ . The Volterra principle continues to hold even in this case, where much of the original model's behavior has changed.

We can continue making the model more complicated by, for example, adding a term for predator satiation. This can be accomplished by incorporating a parameter  $c$  corresponding to the maximum rate of prey capture per predator. Incorporating this term into the density-dependent model, we get

$$\frac{dV}{dt} = r \left( 1 - \frac{V}{K} \right) V - c(1 - e^{-aV/c})P, \quad (9)$$

$$\frac{dP}{dt} = bc(1 - e^{-aV/c})P - mP. \quad (10)$$

Analyzing this model is much more complicated because of the exponential terms and because the model exhibits limit cycles (May 2001). However, upon analyzing the model, we find that the Volterra principle still obtains. The introduction of a general pesticide will increase the relative number of prey.

If we continue this analysis with other models, we will see that the Volterra principle describes a property true of many distinct predator-prey models. This principle is an example of what I call a *robust property*, a dynamic or static property common to many models making different idealizing assumptions.

Discovering robust properties is not the end of robustness analysis. Such a discovery is followed by an investigation of common features of the models to determine which common feature, if any, gives rise to the

property in each case. The structure common to the models that gives rise to the robust property is what I will call the *common structure*. Further analysis of predator-prey models, for example, shows that the Volterra principle arises in “any model in which the abundance of predators is controlled mostly by the growth rate of the prey and the abundance of prey by the death rate of the predators” (Roughgarden 1979, 439).

Putting the common structure and robust property together to form a conditional hypothesis, we can formulate the robust theorem as follows:

Ceteris paribus, if the abundance of predators is controlled mostly by the growth rate of the prey and the abundance of the prey controlled mostly by the death rate of predators, then a general pesticide will increase the abundance of the prey and decrease the abundance of predators.

This robust theorem is a hypothesis about the properties of a large set of actual and possible predator-prey systems, positing the effect of a general pesticide on these systems. It is highly abstract; no mention is made of particular predator-prey systems such as the lynx-hare system, the shark-cod system, or the HIV virus/host cell system. The hypothesis simply describes the effect of a causal structure on a highly general class of systems.

*3.2. Robust Theorems and Robustness Analysis.* The analysis of predator-prey models and the discovery of the Volterra principle provide an excellent template for a more general characterization of robustness analysis and the structure of robust theorems. Robustness analysis can be described as a four-step procedure, although some of the steps can be carried out simultaneously. It begins by examining a group of models to determine if they all predict a common result, the robust property. The second step involves analyzing the models for the common structure that generates the robust property. Results from the first and second steps are combined to formulate the robust theorem itself, a conditional statement linking common structure to robust property, prefaced by a ceteris paribus clause. The third step gives an empirical interpretation of the mathematical structures combined into the conditional form. Finally, the theorist can conduct stability analysis of the robust theorem to determine what conditions will defeat the connection between common structure and robust property. Let's consider some of these steps in more detail.

Robustness analysis begins by examining a group of similar, but distinct, models for a robust behavior. During this stage it is important to collect a sufficiently diverse set of models so that the discovery of a robust property does not depend in an arbitrary way on the set of models analyzed. The first step is either followed by or conducted in parallel with the second,



which involves finding the core structure that gives rise to the robust property.

In straightforward cases, the common structure is literally the same mathematical structure in each model. In such cases one can isolate the common structure and, using mathematical analysis, verify the fact that the common structure gives rise to the robust property. However, such a procedure is not always possible because models may be developed in different mathematical frameworks or may represent a similar causal structure in different ways or at different levels of abstraction. Such cases are much harder to describe in general, relying as they do on the theorist's ability to judge relevantly similar structures. In the most rigorous cases, theorists can demonstrate that each token of the common structure gives rise to the robust behavior and that the tokens of the common structure contain important mathematical similarities, not just intuitive qualitative similarities. However, there are occasions in which theorists rely on judgment and experience, not mathematics or simulation, to make such determinations.

After the first and second stages of robustness analysis, the theorist has formal or mathematical information, but not an empirical description. She knows that one mathematical structure, the common structure, gives rise to another, the robust property. Modeling is typically aimed at telling us something about the real world, and theorists want to know about the properties of real-world phenomena, not mathematical structures. So the third step of robustness analysis involves interpreting the mathematical structures as descriptions of empirical phenomena. In the predation case, theorists have to decide how two coupled differential equations will explicitly map on to the properties of real or imagined predator-prey systems.

After the first three stages of robustness analysis, theorists can formulate robust theorems. They have the following general form:

*Ceteris paribus*, if [common causal structure] obtains, then [robust property] will obtain.

The fourth and final step of robustness analysis consists of various kinds of stability analysis, the purpose of which is to determine what happens to the robust theorem when the situation described by the set of models varies slightly. Consider the predator-prey model example. One way to think about the transition from the model described by equations (2) and (1) to the one described by (7) and (6) is to ask what happens to the behavior of the model when density dependence, even an arbitrarily small amount of it, is factored into the model. Many of the properties of the original model are destroyed, but the Volterra principle remains. So we know that the Volterra principle is insensitive to density dependence. In robustness analysis, this sort of analysis is typically carried out on as

many dimensions as possible. It may ultimately be possible to replace a robust theorem's general *ceteris paribus* clause with a very specific statement of the conditions that defeat the efficacy of the core structure in generating the robust properties.<sup>1</sup>

So far, I have described only the procedures theorists use to discover robust theorems. These procedures are important in themselves, but are not the full story. A common reason theorists engage in robustness analysis is to increase the quality of their predictions and explanations about real phenomena. Although useful for both of these purposes, the theorems generated by robustness analysis cannot fulfill either role alone because they are conditional statements, further attenuated with *ceteris paribus* clauses. Explaining a real-world phenomenon or predicting its occurrence requires us to know that the common structure is actually being instantiated and that no other causal factor is *preempting* the efficacy of the common structure.

One way to determine if the common structure is being instantiated and if any preempting causes are present is to conduct an empirical investigation. While this is the most reliable way to ensure that a robust theorem can be applied, it is often impractical or impossible to collect the relevant data. In fact, robustness analysis is usually introduced in situations in which data are hard to obtain. Fortunately, there is an alternative that, while not completely reliable, can give us good reasons to believe the predictions and explanations of robust theorems. The alternative involves answering two key questions:

1. How frequently is the common structure instantiated in the relevant kind of system?
2. How equal do things have to be in order for the core structure to give rise to the robust property?

Although the first question is best settled empirically, it can be partially addressed using techniques associated with robustness analysis. The key comes in ensuring that a sufficiently heterogeneous set of situations is covered in the set of models subjected to robustness analysis. If a sufficiently heterogeneous set of models for a phenomenon all have the common structure, then it is very likely that the real-world phenomenon has a corresponding causal structure. This would allow us to infer that when we observe the robust property in a real system, then it is likely that the core structure is present and that it is giving rise to the property.

The second question can be more easily addressed as finding that its answer is part of the fourth step of robustness analysis itself. In order to

1. Such a procedure is carried out in Skyrms (2000).

determine how sensitive a robust property is to perturbations, theorists engage in various kinds of stability analyses. If fully carried out, the fourth step of robustness analysis provides enough information to determine what kinds of perturbations will preempt the occurrence of the robust property, even when the core structure is instantiated.

**4. Robustness Analysis's Confirmatory Power.** Orzack and Sober's central criticism of robustness analysis is that it is a nonempirical form of confirmation. Since it is a set of procedures conducted on a set of models, it should be able to generate only mathematical, not empirical, results. Yet the third step of robustness analysis involves generating hypotheses about actual and physically possible empirical phenomena. Is this move from a mathematical fact to an empirical claim as illicit as Orzack and Sober claim?

While the transition from mathematical to empirical may look illicit when described as "nonempirical confirmation," it is actually part of a well-accepted theoretical practice that is so common, it is rarely discussed explicitly. In every scientific domain, theorists must establish that the mathematical framework in which their theories are framed can adequately represent the phenomena of interest. By way of example, consider models of population growth. Standard issues in confirmation theory concern whether a particular kind of model, such as the logistic growth model, is confirmed by the available data. However, there is a prior confirmation-theoretic question that is often asked only implicitly: If the population *is* growing logistically, can the mathematics of the logistic growth model adequately represent this growth? Theorists rarely articulate such questions in research articles, but an affirmative answer underlies their research.

Despite rarely being discussed explicitly, theorists' confidence in their ability to represent phenomena with their models did not come for free. It was minimally established by demonstrating that the relevant mathematics could be deployed to make correct predictions. It may also have been investigated explicitly by mathematicians. These investigations result in what I will call *low-level confirmation*, confirmation of the fact that certain mathematical structures can adequately represent properties of target phenomena.

Low-level confirmation is what allows robust theorems to make claims about real-world phenomena. To see how this works, consider the logistic model once again. If a biologist observes a population for 20 generations and determines that it is growing logistically with a carrying capacity  $K$ , it is reasonable for her to continue using the logistic model to project the growth of the population into the future. This is an empirically risky, but reasonable, hypothesis.

There is an even less risky hypothesis that can be adopted:

If the population is growing logistically, then the logistic model will continue to make accurate predictions about the population's size in the future.

While she cannot know for sure whether the population will continue to grow logistically, this conditional hypothesis seems almost assured because we believe that logistic growth can be accurately represented with the mathematics of the logistic model. Our confidence comes from low-level confirmation.

Robustness analysis also relies on low-level confirmation, insofar as it can produce robust theorems that are confirmed. The first two steps of robustness analysis are concerned with finding the antecedent and consequent that replace the dummy letters in the conditional statement: "Ceteris paribus, if structure  $S$  is found, then property  $P$  will be instantiated." These steps, however, are only mathematical. They help us discover the relationship between  $S$  and  $P$  and may even help us verify this relationship, but this is strictly mathematical; nothing empirical has been demonstrated. Low-level confirmation starts to play a role in the third stage of robustness analysis; it licenses us to regard the mathematical dependence of  $P$  on  $S$  as a causal dependence.

In the predation case, for example, we are confident that ecological relationships can be represented with the models described by coupled differential equations. Thus when we discover the consequences of these models, we are confident that most of these consequences are true of any system described by the model. This confidence comes from low-level confirmation, not from robustness analysis itself. Thus robustness analysis is not a nonempirical form of confirmation as Orzack and Sober suggest. It does not confirm robust theorems; it identifies hypotheses whose confirmation derives from the low-level confirmation of the mathematical framework in which they are embedded.

**5. Conclusions.** Orzack and Sober criticize a simplified version of robustness analysis that involves the identification of robust properties and the treatment of such properties as true descriptions of real phenomena. The analysis presented in this paper shows that the discovery of robust properties is only one part of robustness analysis. A fully formed robust theorem has three parts: a common structure, a robust property, and a set of ceteris paribus conditions.

To use robust theorems in the ways that Orzack and Sober envision, a theorist must take two additional steps beyond discovering a robust theorem. For each phenomenon of interest, the theorist has to argue for the occurrence of the common structure. This could be accomplished

either by direct empirical determination or by gaining evidence about the ubiquity of the structure across similar phenomena of interest. She then must show that there are no preempting causes or other reasons that interfere with the common structure giving rise to the robust property. Only after completing these two steps can a theorist assert, with any certainty, that the robust property described in her robust theorem is being instantiated.

Regarding robustness analysis as a nonempirical form of confirmation is also an oversimplification, one that distorts the origin of the confirmatory weight attached to robust theorems. We are now in a position to see why: Robustness analysis helps to identify robust theorems, but it does not confirm them. Such theorems are confirmed via low-level confirmation, the sort of confirmation that licenses the use of a framework to construct models of phenomena in the first place.

Philosophers of science have begun to recognize the unique features of the sciences, such as population biology, that deal with highly complex phenomena. One of the most important contributions philosophy of science can make is to understand the theoretical tools that must be employed when dealing with such phenomena. Robustness analysis is an indispensable procedure in the arsenal of theorists studying complex phenomena. We must understand it and make room for it in our accounts of theory construction and testing.

## REFERENCES

- Hofbauer, J., and K. Sigmund (1998), *Evolutionary Games and Population Dynamics*. Cambridge: Cambridge University Press.
- Levins, R. (1966), "The Strategy of Model Building in Population Biology," in E. Sober (ed.), *Conceptual Issues in Evolutionary Biology*. 1st ed. Cambridge, MA: MIT Press, 18–27.
- (1993), "A Response to Orzack and Sober: Formal Analysis and the Fluidity of Science," *Quarterly Review of Biology* 68(4): 547–555.
- Lotka, A. J. (1956), *Elements of Mathematical Biology*. New York: Dover.
- May, R. M. (2001), *Stability and Complexity in Model Ecosystems*. Landmarks in Biology ed. Princeton, NJ: Princeton University Press.
- 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(4): 533–546.
- Roughgarden, J. (1979), *Theory of Population Genetics and Evolutionary Ecology: An Introduction*. New York: Macmillan.
- Skyrms, B. (2000), "Stability and Explanatory Significance of Some Evolutionary Models," *Philosophy of Science* 64: 94–113.
- 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.
- 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.