True Lies: Realism, Robustness, and Models

Jay Odenbaugh[†‡]

In this essay, I argue that uneliminated idealizations pose a serious problem for scientific realism. I consider one method for "de-idealizing" models—robustness analysis. However, I argue that unless idealizations are eliminated from an idealized theory and robustness analysis need not do that, scientists are not justified in believing that the theory is true. I consider one example of modeling from the biological sciences that exemplifies the problem.

1. Introduction. In this essay, I first present a problem idealizations pose for scientific realism. Second, I present one way of possibly solving this problem with the technique of robustness analysis. Third, I argue that uneliminated idealizations are a problem for scientific realism, robustness analysis notwithstanding. Fourth, I provide an example of a robustness analysis and uneliminated idealizations with the work of ecologist Henry Horn. Finally, I consider objections to the analysis.

2. The Problem. *Scientific realism* is the claim that scientists are justified in believing that their theories are true, and *antirealism* denies this—they are not justified in believing that their theories are true.[1] Some philosophers of science prefer to define 'scientific realism' as the claim that the aim of science is truth and 'antirealism' as the claim that the aim of science

†To contact the author, please write to: Department of Philosophy, Lewis and Clark College, Portland, OR 97219; e-mail: jay@lclark.edu.

‡I would like to thank Rebecca Copenhaver, Eddie Cushman, Ronald Giere, Wendy Parker, Nicholas D. Smith, and Michael Weisberg for their questions and comments regarding this article.

1. Note that I did not claim that scientific realism requires that scientists have justified true beliefs in their theories since that is tantamount to assuming scientists know (Gettier problems to the side) their theories are true, and that is far too strong.

Philosophy of Science, 78 (December 2011) pp. 1177–1188. 0031-8248/2011/7805-0038$10.00

is empirical adequacy (van Fraassen 1980). However, scientific realism defined axiologically could be true even if scientists never came close to satisfying their aim. Surely scientific realism requires more than an unfulfilled but sought after goal. A *theory* is a set of deductively closed propositions that explain and predict empirical phenomena, and a *model* is a theory that is idealized. Finally, an *idealization* is a false proposition that is useful for the purposes of science qua science.[2]

Confirmed predictions of a theory give one reason to believe that the theory is true.[3] That is, if a theory correctly predicts some phenomenon, then ceteris paribus this is evidence that the theory is true. Put very simply, the problem with which this article is concerned just is this. Consider a model that has a confirmed prediction. First, we know that our model is false. Hence, the confirmed prediction should not increase our confidence in our model. Second, without the idealizations, we have no reason to believe that the nonidealized components of the model imply the confirmed prediction. Hence, the confirmed prediction should not increase our confidence in our model. Probabilistically put, we would say ceteris paribus that the confirmed prediction increases the probability of the theory. However, models pose two problems. First, our theory contains propositions that are false, and hence the probability of the theory is zero. Second, we can characterize the likelihood of the confirmed prediction on the nonidealized components only via the idealizations. Hence, the likelihood of the confirmed prediction on the nonidealized component of the model is undefined.[4] One proposed answer to this problem is that the idealizations are harmless since we can discharge them—our confirmed predictions are *robust*. In this essay, I will examine the property of *robustness* with regard to models to see whether it can be used to resolve this problem.

3. Robustness. Here is how philosopher Bill Wimsatt describes robustness analysis: "The family of criteria and procedures which I seek to describe in their various uses might be called *robustness analysis*. They all involve the following procedures: 1. To analyze a *variety* of *independent* derivation, identification, or measurement processes. 2. To look for and analyze things

2. The purposes of science or scientific theorizing are at least in part explanation, prediction, and intervention. These goals may be considered essential or contingent to the scientific enterprise.

3. Of course, philosophers of science will not simply assert that any confirmed prediction is evidence that a theory is true. The problem will still be present regardless of how one fine-tunes the sort of prediction that is confirmatory for a theory.

4. The probabilistic problem still exists for Bayesians who reject belief and acceptance of scientific theories.

which are *invariant* over or *identical* in the conclusions or results of these processes. 3. To determine the *scope* of the processes across which they are invariant and the *conditions* on which their invariance depends. 4. To analyze and explain any relevant *failures of invariance*" (1981/2007, 2). In the context of mathematical modeling more narrowly, it is customary to consider the work of Richard Levins as providing an explicit discussion of the matter.[5] Here is what Levins says about the notion of *robustness*: "However, even the most flexible models have artificial assumptions. There is always room for doubt as to whether a result depends on the essentials of a model or on the details of the simplifying assumptions. . . . Therefore, we attempt to treat the same problem with several alternative models each with different simplifications but with a common biological assumption. Then, if these models, despite their different assumptions, lead to similar results we have what we can call a robust theorem which is relatively free of the details of the model. Hence our truth is the intersection of independent lies" (1966, 423).[6] As Levins notes, mathematical models in the biological sciences are highly idealized. For any prediction of a model, we can ask (and rightfully so) the following question: "Why believe a model even when its predictions are confirmed?" To see how robustness works, let us consider a historical example regarding evolutionary theory.

During what historians term the "eclipse of Darwinism," antiselectionists argued that natural selection simply could not produce substantial evolutionary change; at best it could "weed out" selectively disadvantageous traits, and something else such as inheritance of acquired characteristics, orthogenesis, or saltationism was required for substantial evolutionary change (Bowler 1983). The eminent biologist and statistician Sir Ronald Fisher articulated models of natural selection in which very small differences in fitness over generational time could take traits to fixation, which was contrary to the suggestions of the antiselectionists (Fisher 1930). So, "Is it possible for small differences in fitness to drive traits to fixation?" is answered yes. However, Fisher's sketched model assumed the population to be infinite in size. Given that idealization, why should the anti-Darwinians have been convinced by Fisher's argument? That is, why should they have believed claims that depended on a false assumption

5. Although Levins is considered one of the first scientists to articulate the notion of robustness, there are a variety of other important scientists including Donald Campbell who have done so as well. In the philosophical literature, Bill Wimsatt has been a tireless defender of the importance of robustness. More recently, Michael Weisberg has provided important analyses of the notion as well (2006).

6. Strictly speaking, one might read Levins as concerned with why one should believe a prediction when it is derived from an idealized model. In this essay, I am concerned with the inverse issue—why believe a model when its predictions are confirmed? Robustness broadly construed is relevant to both issues.

regarding population size? Fisher provides the answer—the claims did not depend on a false assumption. It could be shown that there was no appreciable difference to trait fixation between an infinitely large and a large but finite population. In other words, Fisher showed that trait fixation by natural selection would occur relatively quickly and was robust over different assumptions about population size. In so doing, Fisher's work along with that of Sewall Wright and J. B. Haldane provided "a clearing of the ground of the debris of anti-Darwinian criticism" (289–90).[7]

In order to analyze our problem, I will introduce some technical terms for expository purposes. Let us suppose that theories, and thus models, are described by deductively closed sets of propositions.[8] Each model M_i of a set of models **M** has its propositions divided up into two nonempty subsets. First, there are the *shared assumptions* A that are retained over each element in **M**. Second, there is the complement of A, A^C, that is subdivided into *true* and the *false* propositions, and the idealized assumptions are counted among the latter. Let us call our *prediction* P and suppose it concerns the values or configuration of the variables and parameters. Finally, let us say that two models with A are distinct just in case they contain logically nonequivalent A_i^C and A_j^C assumptions. We can now restate our problem. Suppose the following case: (*a*) P is confirmed, (*b*) (A and A_i^C) entails P (but A does not entail P), and (*c*) A_i^C is idealized. Why believe M since we know that it has an idealized A_i^C and that without it, A would not entail P? Because P does not depend on A_i^C; P is robust.

We can now define the concept of *robustness* (R). Consider a set of models $\mathbf{M} = \{M_1, M_2, \ldots, M_n\}$. Each model is composed of a common A and at least one distinct A_i^C.

> (R) A prediction P is robust over **M** if for each $M_i \in \mathbf{M}$, M_i entails P.

We are now in a position to articulate the concept of a *robustness analysis*. It consists of three steps: articulate a core of assumptions A, devise a cluster of models **M** that retain the core and vary the A_i^C, and for a given

7. I am indebted to Robert Skipper for a discussion of Fisher's accomplishments. The history of this episode is far more complicated than what I have discussed here.

8. I construe theories and models propositionally for ease of exposition, but one could use the semantic view of theories. For example, for some idealized theory, suppose that there is a model that makes the sentences of the theory true and that there is a model that makes the sentences characterizing the confirmed predictions true. Does the fact that the latter model is embeddable in the former give us reason to believe that the idealized parts of the theory are isomorphic to the empirical system?

M, find any robust *P*s (and determine over **M***—where **M** $\subset$ **M***—where *P*s are fragile).[9]

The analysis offered in this article involves some simplifications; however, none of these make things harder for the scientific realist. In fact, some make the case for realism easier rather than harder. For example, we are considering only one idealization per model, the relationship between model and prediction is deductive, we are suppressing auxiliary hypotheses, and the idealizations are permuted one at a time. In the next section, I want to consider whether robustness analysis can provide aid in the defense of scientific realism against idealizations.

4. Realism and Robustness. We can put the issue for the realist as follows: "If *P* is robust over **M**, then idealizations are 'discharged'; however, are we justified in believing *A*?" Here is an example of an argument for suggesting that even after an idealization has been discharged, we still do not have a good reason to believe *P*. Consider the following passage from Orzack and Sober: "When [Levins] further writes that a particular 'nonrobust' 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" (1993, 538). They go on to argue that robustness analysis cannot provide evidence for *P*'s truth independent of empirical confirmation. One cannot have evidence that a contingent claim—a prediction *P*—is true simply from the fact that it follows from a set of models. Of course, this is correct; nevertheless, Levins provides his own response to their argument. He writes, "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" (1993, 554). Hence, he is not suggesting that robustness is a "nonempirical" form of confirmation of *P*. Rather, we have empirical evidence for *A* of **M**, and we have empirical evidence for *P* independent of *A*. However, if we have independent evidence for *A*, then the idealizations in one way do not matter. So, I am concerned with those cases where we do not have empirical evidence for *A* independent of evidence for *P*.

When we do not have independent evidence for *A* and have the problem as I have articulated it: robustness analysis is not sufficient to discharge

9. First-order logic is *monotonic* in the sense that adding propositions to a set does not reduce the former's implications. Is robustness "nonmonotonic"? No. First, note that *A* does not entail *P*. Second, it possible for there to be models (A and A_i^C) and (A and A_j^C) such that the former entails *P* and the latter does not. It is with regard to those latter models that *P* may be "fragile." So, the models in **M** each entail *P*, but there are models in **M*** that do not. Thanks to Eddie Cushman for discussion of this point.

worries about our idealizations. Consider the following argument. Suppose for a prediction P and idealization A_i^C such that $(A \wedge A_i^C)$ entails P, we can provide another A_j^C such that $(C \wedge A_j^C)$ entails P. Thus, we can discharge our worry about A_i^C with A_j^C. However, either A_j^C is an idealization or not. If it is idealization, then we must find some other assumption A_k^C such that $(C \wedge A_k^C)$ entails P. This must continue ad infinitum unless there is at least one A^C that is true. Thus, to discharge our skepticism about our idealizations, a robustness analysis must terminate in an A_j^C that is true (or we have evidence that it is true). But, if we cannot discharge our skepticism regarding our idealizations, then either we do not have a reason to believe A or a reason to believe P can confirm A. Supposing that true A_j^C are not available, then we do not have a reason to believe A or a reason to believe P can confirm A.

In summary, finding invariant predictions over a set of models does not solve our problem for scientific realism. Rather, robustness analysis requires that we have empirical evidence for P independent of A and that we have empirical evidence that at least one A^C is true. Let me now turn to a case study in robustness analysis—the work of Henry Horn on forest succession as a Markov process.

5. Robustness Analysis: Markov Models of Forest Succession. Ecologists are interested in the changes that occur in ecosystems where populations replace each other: *succession*. Succession begins when colonizers arrive in an area and ends when a final, relatively stable state called a 'climax' occurs. In primary succession, colonization occurs where no community is present. In secondary succession, there is an alteration of an already existing community after a disturbance.

Forests change as the result of perturbations like wildfires. Ecologists have noted there are patterns regarding these changes. Similar initial communities follow similar successional stages. Dissimilar pioneer communities arrive at similar final states that resemble stands of virgin forest. Ecologist Henry Horn writes, "The most dramatic property of succession is its repeatable convergence on the same climax community from any of many different starting points. The property is shared by a class of statistical processes known as 'regular Markov chains'" (1975b, 196). In light of this realization, Horn and other ecologists began to model forest succession as Markov chains.

A Markov chain is a stochastic process in which transitions among various states occur with characteristic probabilities that depend only on the current state and not on any previous state. The most important property of regular Markov chains is that they eventually settle into a pattern where the various states occur with characteristic frequencies that are independent of the initial states. The final stationary distribution is thus analogous to the climax community (Horn 1975b, 196).

Horn (1971, 1974, 1975a, 1975b) investigated the forest behind the Institute for Advanced Study in Princeton, New Jersey. These woods have several different stands. One stand was never farmed, and the others were used for various agricultural purposes. However, all had been abandoned for between 30 and 150 years at the time of the study. Horn thus could investigate how they recovered from the "temporary indignities imposed by man or by nature" (1975b, 197) through the process of secondary succession.

Horn supposes that a forest is a honeycomb of independent cells, where a cell can be occupied by one and only one tree. Each cell is roughly the size of a mature tree's crown. Likewise, trees are replaced synchronously by a new generation of saplings that arise from their understory. The probability that a given species will be replaced by another is proportional to the number of saplings of the latter in the understory of the former. Thus, we can think of succession as consisting in a tree-by-tree replacement process. We can estimate the probability that a tree of one species will be replaced by another tree of some species. From a matrix of these probabilities, we can calculate how many trees of each species will be found in any stage of succession.

How exactly do we determine such a matrix of probabilities? Horn, for example, found a total of 837 saplings underneath gray birches scattered through the institute woods. There were no gray birch saplings, 142 red maples, 25 beeches, and so on, in the understory. Therefore, the probability that a gray birch will be replaced by another gray birch is $\Pr(\mathrm{GB/GB}) = 0/837 = 0$. The probability that a gray birch will be replaced by a red maple is $\Pr(\mathrm{RM/GB}) = 142/837 = 0.17$. The probability that a gray birch will be replaced by a beech is $\Pr(\mathrm{B/GB}) = 25/837 = 0.03$.

We can determine the number of gray birches in the next generation by finding all the species in the current canopy that have gray birch saplings in their understory and by multiplying their current abundances by the probability that they will be replaced by a gray birch. So, we have the following equation:

$$\begin{aligned}\mathrm{GB}(t+1) &= \Pr\left(\frac{\mathrm{GB}}{\mathrm{BTA}}\right)\mathrm{BTA}(t) + \Pr\left(\frac{\mathrm{GB}}{\mathrm{SF}}\right)\mathrm{SF}(t) + \Pr\left(\frac{\mathrm{GB}}{\mathrm{BG}}\right)\mathrm{BG}(t)\\ &= 0.05\mathrm{BTA}(t) + 0.01\mathrm{SF}(t) + 0.01\mathrm{BG}(t),\end{aligned}$$

where BTA is a big tooth aspen, SF is sassafras, and BG is black gum.

Generalizing, we can formulate the following discrete model. Let $N_j(t)$ be the proportion of species j in generation t and p_{ij} be the probability that an individual of species j replaces a given individual of species i. Let s be the number of species. So, we have

$$N_j(t+1) = \sum_{i=1}^{s} N_i(t)p_{ij},$$

or in matrix notation with $\mathbf{n}(t)$, a row vector of N_j, and $\mathbf{P}$ and an $s \times s$ matrix of p_{ij}, we have $\mathbf{n}(t+1) = \mathbf{n}(t)\mathbf{P}$. After m generations, we have $\mathbf{n}(t+1) = \mathbf{n}(t)\mathbf{P}^m$. As m gets large, $\mathbf{n}$ will "settle down" to a stationary distribution $\mathbf{n}^*$, which is the solution of s linear equations, $\mathbf{n}^* = \mathbf{n}^*\mathbf{P}$.

On the basis of this model, Horn determined the expected stationary distribution and hence what the climax community should be in the institute woods. He derived the following model:

$$\mathbf{n}^*_{\text{expected}} = \langle 0, 0, 4, 5, 5, 6, 7, 16, 50 \rangle,$$

which is a vector of the stationary distribution percentages. Of course, one of Horn's assumptions in his model is false; namely, trees do not replace each other synchronously. If it were true, then trees must die and be replaced at the very same time. Thus, the discrete model cannot accurately represent the actual course of forest succession. However, the model and the above vector do tell us what the number of occurrences of each species over time should be in a given hypothetical cell. Likewise, if we take a synchronous sample of many cells, we should encounter each species in proportion to the number of times it occurs in the temporal sequence for each cell, weighted by the life span of the species (the average of the life span of each tree of the species; Horn 1975b, 199). Thus, in order to determine the expected abundances in the stationary distribution, we must multiply the above vector by the longevity vector

$$\mathbf{l} = \langle 80, 50, 100, 150, 200, 300, 200, 250, 200, 15, 30 \rangle$$

and normalize the products. So, the age-corrected stationary distribution is

$$\mathbf{n}^*_{\text{expected}} = \langle 0, 0, 2, 3, 4, 3, 4, 6, 6, 10, 63 \rangle.$$

Finally, the observed vector is

$$\mathbf{n}^*_{\text{expected}} = \langle 0, 0, 0, 6, 0, 3, 0, 0, 14, 1, 76 \rangle.$$

Horn did not attempt to estimate the goodness of fit of $\mathbf{n}^*_{\text{expected}}$ to $\mathbf{n}^*_{\text{observed}}$. Nonetheless, it is reasonable to conclude that Horn's analysis of the Princeton University woods bears a good fit, albeit rough, to the phenomena. However, Horn does not claim that his model accurately represents the actual course of succession but only that "the stationary distribution, when weighted by the longevity of each species, should represent the actual distribution of species in the climax" (1975b, 200). Hence, it is claimed only that the expected and the observed climax community are in rough agreement.

What about the false assumptions of Horn's model? Horn writes, "I shall routinely make outrageous assumptions, but I shall defend them in several ways. Some are needed only for analytic convenience and may be relaxed with no major effect on the result. In some cases a redefinition of a measurement is all that is needed to bring theory into line with fact. Astoundingly, some of the assumptions are even true" (1975b, 197). In the case of the synchrony assumption, Horn argues that the assumption does not matter—he discharges the assumption in two ways. First, Horn shows that the idealization of synchrony can be discharged by correcting it with respect to the phenomena (through weighting the relative number of occurrences of each species by their longevity). This assumption is used simply because it allows for mathematical ease. Second, Horn also provides models that do not assume synchrony as well (206–8). He can show that some result does not uniquely depend on the idealized assumption through a robustness analysis, although he does not call it that by name. Horn implicitly uses a robustness analysis when he claims that "all the properties of the rudimentary model are shared by a more realistic model which allows overlapping generations and diverse rates of survival for different species at each of several stages in their life histories" (198).

What is the more realistic model? Let each species of tree j have a characteristic and constant mortality rate d_j. Thus, the rate of increase dN_j/dt of each species j has a term $-d_jN_j$ that is due to deaths and a term that sums the number of deaths of other trees multiplied by the probability that species j will replace the other dying trees. So, we have

$$\frac{dN_j}{dt} = -d_jN_j + \sum_i d_iN_ip_{ij}.$$

Thus, Horn builds asynchrony and varied life spans directly into the model, as opposed to correcting for it with respect to the phenomena. By replacing the synchrony assumptions with the asynchrony and varied life span, what are the relevant robust predictions? The relevant prediction is the stable stationary distribution, or "climax community," which is a prediction from both models (1975a). Horn also reminds us why we do not always try to build more realistic models: "I can and shall add varied life spans and synchrony to the next model, but I shall leave the fiendish empirical computations of such a model for a later paper" (1975b, 201). Incidentally, this paper with the fiendish computations never appeared.

Let us suppose that Horn successfully discharges the idealization that "trees die and are replaced at the very same time." Even so, he does not conduct a robustness analysis regarding other putative idealizations: the matrix of transition probabilities is homogenous, there are no time lags, and there is no density dependence. If he has not shown these idealizations are eliminable, then his Markov model (or some component of it) should not be believed to be true. Moreover, confirmed predictions would not confirm the model (or some component of it) either. That is, he should not say, "The most dramatic property of succession is its repeatable convergence on the same climax community from any of many different starting points. The property is shared by a class of statistical processes known as 'regular Markov chains'" (1975b, 196). The fact that Horn's model has confirmed predictions does not provide good reason for believing that secondary succession is a Markov process, given that he has not eliminated the many idealizations of his model.

Although I have just considered one example from community ecology, I would suggest that very often in the sciences idealizations are not eliminated, and when they are not, scientists are unjustified in believing that their theories are true. Moreover, robustness analysis can aid them in eliminating idealizations, but more is required.

6. Objections. I now consider objections.

Objection.—Often we are worried about specific idealizations and not idealizations per se. But, by my own admission, robustness analysis can discharge those specific idealizations. This objection depends on a kind of scientific "contextualism" analogous to that in epistemology—a possibility may be ignored in one context but not in another. In response, we worry about idealizations because they are false assumptions. Hence, I cannot see how one could worry about one idealization because it is false without worrying about the others since our reason for worrying about the former applies to the latter with equal force.[10]

Objection.—"Metaphysical realism" concerning truth and falsity of theories is irrelevant to the sciences. Rather, we should be considering how similar models and empirical systems are or whether the assumptions of the models are "approximately true."[11] In response, first, the issues can be raised within a deflationary theory of truth; it need not be "metaphysical." Second, given the absence of independent evidence for the similarity between the model core and the system, we must infer it from the similarity of the predictions and the observations. Given the known dissimilarity between the former, we do not have a reason to believe this helps. Third, for all the work done on approximate truth, there is simply

10. Sometimes false assumptions are introduced via the mathematics employed. Scientists may bracket these because they concern matters more remote for their day-to-day concerns. However, on some occasions these mathematical niceties intersect issues of relevance (e.g., a population growing discretely may not behave the same as one growing continuously).

11. In discussion at the PSA 2010, Ronald Giere raised this objection or something quite similar.

no approach we can import into the problem. And, even if we could, we still know that the core and the system are very, very different. How does the assumption that all trees in a forest die and the new trees are born at the same time approximate the assumption that their generations overlap with birth and death occurring all the while? If two propositions are logically inconsistent, then in what sense does the former approximate the latter?[12]

Objection.—If a model is more predictively accurate than another (e.g., the likelihood of the former on the data is greater than the latter on the data), then surely this is evidence that the former is less idealized than the latter. A simple example suffices to challenge this claim. Suppose I am in Oregon and you want to know what sort of weather I am experiencing. One hypothesis is that I am in the United States, and the other is that I am in Washington state. The former is true and the latter false, but the latter is more predictively accurate of the weather that I experience. Truth need not be more accurate than falsity.

In this essay, I have not been sketching a global, or a priori, argument for antirealism. In fact, nothing I have said is contrary to scientific realism when idealizations are not present or when they can be eliminated. Rather, I have been solely concerned with those circumstances when they have not been eliminated, which I believe to be more frequent than one might assume. The work of Horn in community ecology was meant as an example of this.

7. Conclusion. I have argued that one is justified in believing a theory only if its idealizations are eliminable, and robustness analysis may not do this. Moreover, I have provided one example, that of Henry Horn's Markov models of forest succession. Although he successfully eliminated at least one of his idealizations, the confirmation of his predictions was not sufficient for believing that forest succession is a Markov process. More generally, if idealizations are generally ineliminable, we are rarely justified in believing our models. Our truth is nowhere in the lies.

REFERENCES

Bowler, P. 1983. *The Eclipse of Darwinism: Anti-Darwinian Evolutionary Theories in the Decades around 1900*. Baltimore: Johns Hopkins University Press.
Fisher, R. A. 1930. *The Genetical Theory of Natural Selection*. Oxford: Clarendon.
Horn, H. 1971. *The Adaptive Geometry of Trees*. Princeton, NJ: Princeton University Press.
———. 1974. "The Ecology of Secondary Succession." *Annual Review of Ecology and Systematics* 5:25–37.

12. In simple cases, this might make sense. For example, we know exponential and logistic growth are approximately the same when $N/K \ll 1$ since $dN/dt = rN(1 - N/K) \approx rN$. However, this works only when one model is a limiting case of another, which is exceptional.

———. 1975a. "Forest Succession." *Scientific American* 232:90–98.

———. 1975b. "Markovian Properties of Forest Succession." In *Ecology and Evolution of Communities*, ed. M. L. Cody and J. M. Diamond. Cambridge, MA: Harvard University Press.

Levins, R. 1966. "The Strategy of Model Building in Population Biology." *American Scientist* 54:421–31.

———. 1993. "A Response to Orzack and Sober: Formal Analysis and the Fluidity of Science." *Quarterly Review of Biology* 68:547–55.

Orzack, S., and E. Sober 1993. "A Critical Assessment of Levins' 'The Strategy of Model Building' (1966)." *Quarterly Review of Biology* 68:534–46.

van Fraassen, B. 1980. *The Scientific Image*. Oxford: Oxford University Press.

Weisberg, M. 2006. "Robustness Analysis." *Philosophy of Science* 73:730–42.

Wimsatt, W. 1981/2007. "Robustness, Reliability, and Overdetermination." In *Scientific Inquiry in the Social Sciences: A Volume in Honor of Donald T. Campbell*, ed. M. Brewer and B. Collins, 123–62. San Francisco: Jossey-Bass. Repr. in *Re-engineering Philosophy for Limited Beings: Piecewise Approximations to Reality* (Cambridge, MA: Harvard University Press, 2007).