

Forty Years of ‘The Strategy’

Levins on Model Building and Idealization

Michael Weisberg (weisberg@phil.upenn.edu)

University of Pennsylvania

Abstract.

This paper is an interpretation and defense of Richard Levins’ “The Strategy of Model Building in Population Biology.” Levins’ “The Strategy” has been extremely influential among biologists since its publication forty years ago. In this influential article, Levins confronts some of the deepest philosophical issues surrounding modeling and theory construction. By way of interpretation, I will discuss each of Levins’ major philosophical themes: the problem of complexity, the existence and consequence of tradeoffs, and robustness analysis. I will argue that Levins’ article is concerned, at its core, with justifying the use of multiple, idealized models in population biology.

Keywords: Richard Levins, tradeoffs, models, model building, idealization, robustness, population biology, theoretical ecology, complex systems

Richard Levins’ article “The Strategy of Model Building in Population Biology” is one of the most influential philosophical discussions about theory construction. It is frequently cited by population biologists and while less known among philosophers of science, it was included in the first edition of Elliott Sober’s widely used philosophy of biology anthology. (Sober, 1984)

Levins’ article contains enormous insight into some of the deepest philosophical issues surrounding modeling and theory construction. His focus on the practice of modeling, rather than on the outcome of mature theorizing, revealed important aspects of the enterprise that had not previously been appreciated. For example, he called attention to the tradeoffs one faces when constructing and analyzing models. He also thoughtfully discussed different strategies for constructing models, the complex ways that one gathers confirming evidence for these models, and the pitfalls associated with aiming for complete representations at the expense of all other considerations.

Despite the importance of Levins’ article, it has received very little critical attention by philosophers of science. Notable exceptions include articles by William Wimsatt (1987, 1981) and a joint article by Stephen Orzack and Elliott Sober (1993). Within the last few years, several discussions of Levins’ article and related methodological discussions have started to appear (Odenbaugh, ms; Weisberg, 2004; Weisberg, forthcoming-a) and a conference has been devoted to “The Strategy of Model Building” and Levins’ methodological work, from which many of the papers in this issue of *Biology and Philosophy* are drawn.

This paper is an interpretation and defense of “The Strategy of Model Building.” I will discuss each of the major themes of the paper: the challenge of complexity, the brute-force approach, the strategies of modeling, and robustness analysis. I will also argue that Levins’ article is concerned, at its core, with the justification of idealization in population biology. The methods developed in this paper can be used to develop a more comprehensive theory.

1. Model Building and its Goals

Levins’ article is about *model building*, which despite its obvious importance in theoretical practice, is rarely discussed in its own right by philosophers. The literature about models has tended to focus on the structure of models, the autonomy of models from theories (Morgan and Morrison, 1999), and the ways that theories can be reconstructed as clusters of models (Suppes, 1960a; Suppes, 1960b; Suppe, 1977). These issues are all relevant to understanding the significance of Levins’ position, but do not concern the primary locus of his discussion: the *practice* of constructing and analyzing models, or what Levins calls model building.

Model building or modeling is the indirect representation and analysis of a real-world phenomenon using a model. It takes place in three stages: In the first, a theorist constructs a model, typically by writing down a mathematical description of this model. In the second, she analyzes the model, looking for characteristic behaviors such as equilibria, oscillations, regions of stability, etc. Finally, if warranted by the problem of interest,¹ the modeler assess the relationship between the model and real-world phenomena. This relationship is assessed using the modeler’s construal, a set of intentions about how the parts of the model should map on to the phenomenon of interest as well as the standards of fidelity that will be used in evaluating the success of the model’s representation. If the model successfully represents a real world phenomenon, then the representation and analysis of the model is, indirectly, a representation and analysis of a real world phenomenon. (Weisberg, forthcoming-b)

Modeling has a venerable history in population biology. Its origin is often traced to the independent, but convergent work of Vito Volterra and Alfred Lotka on predator-prey oscillations. Volterra’s work

¹ I put this point carefully because some models are studied for their intrinsic interest, with no expectation that a real-world phenomenon corresponds to them. Analyses of perpetual motion machine models, three-sex biology models, and models of non-aromatic cyclohexatriene are examples of such a situation.

on predation began as a response to unusual data from the Adriatic fisheries. This data remained mysterious until he constructed a very simple mathematical model of a single population of predators and a single population of prey. After performing a detailed analysis of this model, which apparently replicated the unusual behavior observed in the Adriatic, Volterra argued that the model could approximately replicate the population dynamics of the Adriatic fisheries. (Volterra, 1926) This, along with early work on mathematical population genetics, introduced modeling as a distinct kind of theorizing in population biology. Today, much of biological theorizing involves modeling.

The value of different types of theory construction is a complex and not particularly well-understood aspect of scientific method. Levins, however, offers us three reasons why we should engage in modeling: understanding, predicting, and modifying nature. He calls these the goals of modeling. (Levins, 1966, 19)

A model which is constructed to explain some phenomenon is one that meets the goal of understanding. Similarly, the goal of predicting is met by a model which can make accurate predictions. Modifying nature is a more complex goal because there is no one particular theoretical virtue associated with intervening or modifying nature. In fact, one might think that modifying nature requires both the ability to explain and the ability to predict. So perhaps there are really two goals of modeling or else modifying nature is a goal that requires theoretical desiderata beyond explanatory and predictive power.

Levins' discussion of the three goals of modeling might seem rather commonplace from the point of view of philosophy of science. After all, it is often said that science is about explanation, prediction, and control, and Levins' three goals seem to correspond almost exactly to that maxim. But their simplicity and parallels to standard conceptions in philosophy of science are deceiving. Although Levins initially leaves open the possibility that a single model could achieve these goals, his later discussion emphasizes that the value of these goals are distinct and achieving one of the goals in the maximum degree might be incompatible with achieving another.

2. Complexity and The Brute Force Approach

In my view, the real core of "The Strategy" begins with Levins' discussion of complexity. He reminds us that even relatively simple populations, if represented in complete detail, would require dealing simultaneously with "genetic, physiological, and age heterogeneity within species of multi-species systems that are changing demographically and

evolving under fluctuating influences from the environment.” (Levins, 1966, 18) In other words, even the internal dynamics of a population, completely isolated from external environments, contain an enormous amount of structure to be modeled. Since real populations *are* exposed to environments, this leads to further dynamics in the population including natural selection. In addition, external environments change and are changed by the population over time. Trying to capture all of these factors for a very simple population may be impossible.

When confronted with this much complexity, we have two main options in how we approach theoretical representation. We can adopt what Levins calls the *brute-force* approach, or we can adopt an idealization approach. Either we aim at building as much of the target system’s complexity into our models as we possibly can, or we choose to make *strategic* idealizations, omitting select aspects of the complexity. Much of Levins’ discussion in “The Strategy” and his subsequent philosophical discussions (Levins, 1968; Levins, 1993; Levins and Lewontin, 1985) analyzes the advantages and disadvantages of the idealization approach. Such an analysis most naturally takes place in contrast to the advantages and disadvantages of the brute-force approach.

2.1. THE BRUTE-FORCE APPROACH

Levins describes the brute-force approach as one where we build mathematical models which are “faithful, one-to-one reflection[s] of this complexity.” One could read this in two ways: Is Levins referring to actually constructing models which are one-to-one reflections of the complexity or to the goal of building such models? I think that the most natural way to read this passage is as referring to the goal, not the achievement, complete representation. Elsewhere, I have called the goals and ideals associated with modeling *representational ideals* and named the ideal associated with the brute-force approach COMPLETENESS. (Weisberg, ms-b)

When a modeler adopts COMPLETENESS, she sets her standards of fidelity such that the best model is the one which is a complete representation. Each aspect of the target phenomenon must be mappable onto an aspect of the model. Anything external to the phenomenon that gives rise to its properties must also be included in the representation. Structural and causal connections within the target phenomenon must be reflected in the structure of the representation. Finally, according to COMPLETENESS, the best representation is one which represents all aspects of the target phenomenon with an arbitrarily high degree of precision and accuracy.

Completeness is rarely, if ever achieved, by a model, especially when one is dealing with complex systems. Thus adopting the brute-force approach does not mean that scientists are literally trying to find models that achieve complete representation. Nevertheless, adopting the brute-force approach and its ideal of COMPLETENESS plays two roles in theoretical inquiry. Firstly, it sets up a linear scale along which we can measure the goodness of any model for a given system both in terms of its predictive power and also the causal structure represented by the model. The closer the model comes to complete representation, the better the model is. The second function of the ideal is to set up a target that theorists aim at, but don't literally strive to achieve. As they come closer to the target, theorists can conclude that their models are more successful representations of real-world systems. However, they needn't think that they will ever actually achieve this ideal of complete representation. This kind of ideal is very similar to what Kant calls a *regulative ideal*. It is an ideal which guides inquiry and sets a target to aim at, but the target is known to be impossible to achieve in its entirety.

Levins probably had *systems ecology* in mind when discussing the brute-force approach. Systems ecologists aim to capture as much of the detailed interaction of complex ecosystems as is possible and take it as their continuing mission to include ever more of the world's complexity in their models. Even when pointing to the great difficulty in achieving a complete representation, they make clear how significant their goals actually are.

Simulation models of ecosystems need not be all inclusive but they must be all encompassing, i.e. they must cover all the kinds of interactions present in the system without including all the interactions. (Wiegert, 1975, 314)

They rely on the most detailed descriptions of ecological systems that are currently available. As systems models are always computer-based, very little attention is paid to simplicity, generality, ease of computation, and the like, except insofar as this impacts computational tractability. (Odum, 1983; Watt, 1956) The models are often designed in a modular way, which allow different teams of researchers to contribute different parts of the model. Since they are extremely detailed representations of particular ecosystems, these models are often used to forecast changes in environments or the effect that biotic and abiotic interventions would have on the environment.

In some ways, the brute-force approach is the natural extension of standard, realist accounts of scientific method and theory confirmation. Confirmation theory tries to account for how evidence bears on the

truth or falsity of a theory. As such, philosophers developing these accounts assume that theorists are trying to give complete and accurate representations of the phenomena under consideration. Approximations may be necessary for pragmatic reasons, but this is simply a less-than-desirable state of affairs, something to be eliminated as science progresses. It is certainly not thought to be a conscious choice in theory construction. Levins believes this perspective is mistaken and is highly critical of the brute-force approach when dealing with complex systems. He gives three reasons to be critical of this approach.

2.2. PROBLEMS OF MEASUREMENT

Practical and fundamental limitations on measurement comprise Levins' first critique of the brute-force approach. He writes that brute-force models “[have] too many parameters to measure; some are still only vaguely defined; many would require a lifetime each for their measurement.” (Levins, 1966, 18) Put differently, a major impediment to brute-force modeling is the limitation of our ability to acquire the relevant data for such a model.

This may not be a fully satisfying condemnation of brute-force modeling because it relies on highly contingent facts. However, if it could be shown that making the measurements necessary to build brute-force models was always going to be impossible, this argument would give us an *in principle* reason why the brute-force approach was inadvisable.

In some situations, I think Levins is right that the difficulty of collecting data can undermine the brute-force approach. Consider a phenomenon that behaves in a simple way when looked at in large scale, but that exhibits fluctuations at micro-scales such as the ocean. For example, say we wanted to build a highly accurate model of the coastal waters of California using the brute-force approach. Even though there may be large-scale trends in this systems, there will be an enormous number of micro-fluctuations on small scales. In order to capture all of the relevant detail, we would need to lay down many instrument buoys to record information about water temperature, speed, direction, salinity, etc.

How many buoys should we or can we lay down? In an extremely precise study, maybe they would be spaced apart in one kilometer increments. But let's be unrealistic, say we spaced them one meter apart from one another. Even if we expended enormous resources to get down to this scale, there still would be many fluctuations at even smaller scales that we would not be able to capture. If the instruments are spaced apart in meter increments, we are still 10 orders of magnitude from the molecular level, where many of these micro-fluctuations are

observed. Even worse, putting the measuring devices this close together might have significant effects on the system we were trying to measure. Instruments spaced so closely together would almost certainly affect the flow rate of the water, the local populations of fish, and even possibly the temperature of the water.²

In such a case, even coming close to constructing a brute-force model may be far beyond our ability to carry out. For a relatively large phenomenon, such as the currents in the California coastal waters, a lifetime of work and trillions of dollars wouldn't be enough to alleviate this problem. If you want to look at a large system in microscopic detail, it will be nearly impossible to measure the state of every microscopic component for even a short amount of time. And I only say "nearly impossible" because although far more than merely "practically impossible," it isn't strictly physically impossible.

Levins, of course, has examples from population biology in mind, where the relevant phenomena are macroscopic populations of organisms and the environments in which these organisms are living. He tells us that a brute-force approach to these systems would require measuring hundreds of parameters. Since these parameters will be things like temperature, life-span, availability of food, etc., there won't be as great a measurement problem as in the oceanographic case. However, as any detailed recounting of field studies clearly suggests, field measurement are difficult, tedious, and time consuming. Although in these cases the brute-force approach is in better shape than the oceanography case, gathering the requisite measurements is far from easy and may be outside the realm of possibility even with the largest imaginable group of eager graduate students.

The considerations I have raised so far in connection with Levins' first criticism suggest that it will be impossible to carry out brute-force model building to completion. However, none of the considerations I have raised suggest we couldn't get started on such a program. As I have already explained, building the perfect brute-force model is not a goal any theorist really believes that she can achieve. So the important question for us to ask is whether these considerations undermine the *ideal* of brute-force modeling.

There is not a single answer to this question because it needs to be evaluated with respect to the three different goals of model building, but let's focus on just one of these: building models in order to make accurate predictions. For some kinds of target phenomena, this criticism does not undermine the brute-force approach to constructing predictive models. In particular, the brute-force approach can avoid such

² [Citation to Levins' contribution in this volume]

measurement problems as long as the following are true: 1) The causal forces giving rise to the phenomenon being modeled can be sorted into primary and secondary causes, the primary ones being responsible for most of the behavior; and 2) The same causal forces are acting on all of the parts of the system. If these facts are true of the system being modeled, then the brute-force program can avoid the measurement problems raised by Levins. This is fairly trivial result because these conditions ensure that measurement will be straightforward. Statistical samples of such a system will reflect the behavior of the whole system.

Once we move into the realm of systems where these two conditions are not met, the situation become more complicated. Let's separate those cases where taking measurements actually interferes with the phenomenon we are trying to measure and those cases where it does not. In the first kind of case, the brute-force approach is a self-defeating ideal and hence clearly not a good one. If in order to construct a model, we have to collect data in ways that destroy the very phenomenon we are trying to model, we obviously need to adopt a different strategy.

The much harder type of case is when measurements won't actually interfere with the target system, but Levins' first criticism still obtains. In such a case, it will be practically impossible to carry out the measurements required for a complete representation of the target system or even to come close to carrying it out. Clearly the achievement of complete representation is not a serious possibility in this case. So if the brute-force approach is justifiable, it must be justified on grounds other than the eventual achievement of a complete representation. One way to justify the brute-force approach is if it leads to an active research program, one that will have many useful results at intermediate stages of research. If the approach does have this desirable consequence, then it is a reasonable ideal to have. However, when the brute-force strategy doesn't generate useful results at intermediate stages, it is no longer a strategically sound way to proceed. This determination must be made in particular circumstances; no global verdict is possible.

2.3. ANALYTICAL SOLUTIONS

We now move from an experimental problem with the brute-force approach to the more theoretical ones. The second problem Levins raises points to the difficulty or impossibility involved in solving the equations which describe complex systems. Levins writes, "The equations are insoluble analytically and exceed the capacities of even good computers." (Levins, 1966, 19) Although this statement of the problem is very succinct, there are actually two aspects of the problem that Levins has picked out. The first aspect of the problem is that models of

complex phenomena are not solvable analytically. The second part of the problem is that the complex equations which describe complicated models are often not solvable numerically by computers. This second aspect of the problem is far less significant now than it was in the 1960s when Levins wrote “The Strategy.” Some kinds of brute-force modeling such as long-term weather forecasting still exceeds our computational capacities, but this is a problem that changes with the availability of fast, relatively inexpensive computers.

An analytic solution to a mathematical problem is one “that can be written in ‘closed form’ in terms of known functions, constants, etc . . .” (Weisstein, 2003, 73) In most of the modeling contexts that we are concerned with, the term “analytical solution” refers to a function which is the solution of a set of differential equations which constitute the description of a model. Let us grant Levins the claim that many or most complete models of complex phenomena will be described by equations which are not solvable analytically. How does this pose a problem for the brute-force approach and for which uses of models is this problem most acute?

The lack of analytic solutions does not seem like a problem when we want to use models to make predictions. Say we want to build a very complicated model of ocean currents off the coast of California. If we can perform the relevant calculations, it doesn’t matter whether they are carried out using extremely complex numerical techniques that require several days of computer time or if we can solve the equations analytically. Of course, it would be more convenient to solve the equations analytically in that the computational effort we need to expend would be diminished. However, if our numerical techniques are sound, we wouldn’t get a much better prediction by analytical means.

Constructing models for explanatory purposes may be quite different. Theorists often claim that there is a real advantage to solving equations analytically and hence finding models that allow for an analytical treatment. The advantage is that when one has an analytical solution to an equation, one has an explicit description of how the parts of the model depend on one another and the magnitude of these dependencies. Seeing that fur thickness is inversely proportional to basal core body temperature for different populations of the same species tells you far more about how a system is working than the numerical or graphical data that numerical analysis gives you. When you want to show how and why a system behaves in the way that it does and the system admits of mathematical representation, analytical solutions

can show these dependencies explicitly.³ Hence they are held to be the gold standard by many theorists.

Analytical solutions also have another kind of explanatory virtue. They help us evaluate how general a family of models is. If we simply have numerical solutions to the equations which describe a family of models, then it is very difficult to know how many phenomena can be described by this family. A single numerical solution will not give us a complete characterization of the model because we will have only calculated what will happen with very specific boundary conditions. Many numerical solutions could potentially characterize a single model in its entirety. But if we want to know about the generality of a family of models represented by the same equation, numerical solutions are usually thought not to give us this information.

Although these issues suggest important, non-pragmatic reasons to prefer analytical solutions, I think Levins may be overstating the importance of such solutions. What is important is not the analytical solution per se, but rather having a complete characterization of the dynamics of a model or set of models. Analytical solutions to the equations describing a model will give us an exact characterizations of the entire family of models that this equation could potentially describe. However, some systems of equations can be analyzed using techniques such as local stability analysis. This can give us a relatively full picture of the regions of attraction and repulsion, which points are stable, surfaces of neutral stability, etc. Such an analysis will also allow us to generalize from a single instantiation of a model description to the entire family of models that can be described with the same uninstantiated description.⁴

Of course, some models generated using the brute-force approach will not admit of this kind of analysis; they will be simply too complex for such a complete characterization. So determining the magnitude of this problem would really involve the hard work of looking, model-by-model, to see if there is any pattern in what kinds of models generate mathematically tractable state spaces that can be analyzed and the ones that do not. Insofar as their state spaces can be fully characterized, even without analytical solutions to the equations, we can see how general these models are. If they cannot be analyzed, then Levins' second critique of using brute force models for explanatory purposes stands.

³ I thank Glen Ierley for pointing out why many theorists see this as the main advantage of analytical solutions.

⁴ In conversation, Grigori Mints suggested that such a full analysis gives you essentially everything you would want from an analytical solution. Thus it is not entirely clear that such a full characterization isn't some form of an analytical solution, although not an algebraic solution expressed in closed form.

Given the difficulties in *using* brute-force type models for explanatory purposes, let us now consider whether or not the lack of analytic solutions interferes with *adopting* the representational ideal of COMPLETENESS. Similar to Levins' first criticism of the brute-force approach, there is no single answer to this question; it depends on the kind of phenomenon one is trying to characterize. Take a simple physical system that admits of no analytical solution, such as a three mass system with gravitational attraction between the masses. Although this type of system will admit of no analytic solution in closed form, this model is described by a system of equations that can be analyzed in great detail. We can produce a relatively complete analysis of the state space associated with this system and hence know what kind of behaviors to expect from this and related models. So the lack of analytic solutions does not pose a problem for this and related physical systems, but how about biological ones?

One is tempted to say that this must also be true of many biological systems. The dynamics of two-locus models, for example, cannot be evaluated analytically, but the state space associated with these models can be analyzed in detail. But this is of small comfort. Two-locus models are only one small step toward the ideal of the brute-force research program. Many details of actual genetic systems have been left out of the two-locus model and would also be left out of a ten-locus model. Although it is possible that there is some real genetic system that behaves exactly like the two-locus model, this is generally not the case. So we should be very careful to note that this sort of defense of brute-force modeling as a fruitful goal is only applicable when we are aiming at complete representation.

The more interesting case is where no general analysis of the model can be given at all. In this case, a model generated by the brute-force approach will score poorly on both of the explanatory virtues I have discussed. If a research program centered around generating brute-force models in this domain could hope to achieve these ends, then perhaps we could adopt the ideal of complete representation. However, these are not the sorts of properties that are likely to fall into place with further research. More data or greater computational power simply will not help. If there are no analytical solutions and the dynamics of the model cannot be characterized in some other way, this is a mathematical limitation. Thus a research program which had complete representation as its representational ideal and required the use of enormously complex sets of equations to describe its models would not make progress toward its goal at all. It would be better in such a case to take the construction of idealized models as a goal, perhaps with the proviso that the models should represent the most important causal factors.

2.4. BRUTE-FORCE MODELS WOULD HAVE NO MEANING FOR US

While Levins' first two criticisms of the brute-force approach deal with fundamental issues confronting modelers, it is perhaps his third criticism which has the greatest philosophical insight and significance. Apparently conceding the possibility that numerical solutions for extremely complicated models *may* be possible one day, Levins writes that "Even if soluble, the result expressed in the form of quotients of sums of products of parameters would have no meaning for us." (Levins, 1966, 19) Even if we can perform calculations based on very complicated models using computers, those solutions will not be valuable since they "have no meaning for us." I think that understanding this claim is the key to understanding much of Levins' perspective on modeling. What does it mean to say that a numerical solution has no meaning for us?

I believe that the expression 'meaning for us' means 'understanding the model.' Thus Levins is claiming that while we might be able to make excellent predictions with brute-force models if we can numerically solve the relevant equations using a computer, such complex calculations will leave us completely in the dark about *why* the target system behaves the way it does. It tells us nothing about how the system works and would provide little guidance about how similar systems behave.

Levins' other criticisms of the brute-force approach are largely pragmatic, but the third criticism speaks to a much deeper issue. Levins claims that brute-force generated models are not explanatory. Since almost all scientists and philosophers accept explanation as one of the most important goals of theoretical practice, this criticism effectively undermines the brute-force approach if it is correct.

Unfortunately, Levins does not elaborate on the incompatibility between the brute-force approach and explanation in "The Strategy" or other work. However, there are several possible ways to defend Levins' position. The first possibility points to the cognitive limitations of scientists and other consumers of scientific theories. As Levins indicates at the beginning of his section on the brute-force approach, these models are going to be really complicated, often consisting of hundreds of differential equations. Humans are incapable of grasping such complex dynamics and the associated several hundred dimensional state-space. These models cannot really be thought about as wholes, they can only be manipulated on computers and, perhaps, thought about a bit at a time. Little philosophical work has been produced on this issue, but I think it is an important one to pursue.

Another possibility is that in their attempt to be complete, brute-force models give us little insight about the relative importance of the factors giving rise to a behavior of interest. If I throw a brick through

the window, the complete causal explanation of this event will be very detailed and include information about the molecular structure of the glass, the density of the brick, the brick's exact velocity, its mass, and so forth. Most of these details do not make a difference to the occurrence of the breaking and what does explain the glass being broken is that an object was thrown with such and such a momentum at a fragile window. (Strevens, 2004) A good explanation will discriminate the difference makers from the rest of the causal factors which played a role in generating the explanandum. Brute-force models include all of the factors, but our explanatory practices may demand more discrimination.

Finally, and this may be related to finding the difference makers, brute-force models are tailored to very specific phenomena. They will often not generalize beyond the particular phenomenon one is considering and they will have very limited applicability to other possible, but non-actual phenomena. These two types of generalizability are often thought to be connected to scientific explanation (Weisberg, 2004). If any of the accounts of scientific explanation connecting generality to explanation turn out to be correct, then brute-force models are unlikely to be explanatory.

Thus there are several avenues one might pursue in defending Levins' claim that brute-force models will have no meaning for us. Each of these avenues is worth pursuing for its own sake and for helping us gain greater insight into the problems with the brute-force approach envisioned by Levins.

3. The Idealization Approach

In contrast to adopting the brute-force approach, one might accept from the outset that some parts of the phenomena of interest are not going to be represented in our models. We no longer even *aim* at producing complete models. When we build a model of some phenomenon, we will try to include only the most important or the most relevant aspects, which depends on the interests of the theorist.

A special problem confronts us when we adopt the idealization approach. How do we know when to be satisfied with the accuracy of our models, given that we have committed ourselves to their being inaccurate from the start? In the brute-force approach, there is a simple answer to this sort of question. The aim is complete representation of all of the causal forces that give rise to the behavior of the target system to arbitrary degrees of accuracy and precision. Although almost every real model falls short of this aim, brute-force modelers still believe

that a more complete representation is always a better one. Thus the brute-force approach gives us a highly principled way to structure our modeling practices.

The idealization approach obviously calls for different representational ideals, but which ones? Enumerating the representational ideals associated with the idealization approach is very difficult, for there are indefinitely many ways to idealize, only some of them valuable. Even generating a list of the valuable ones is difficult, but what we really want is an account that not only lists them but that can advise us about which ones to adopt under which kinds of circumstances. Different types of idealizations will result in models with different sets of virtues. No one has ever given such an account, yet “The Strategy” is an important first step as it outlines three approaches to modeling, each reflecting a different representational ideal.

Along with the complication of having many dimensions along which to relax the representational ideal itself, Levins raises another problem in finding the principles to guide idealization: the *historical grounding* of theorizing. During different stages in the development of a science, different idealizations are appropriate. Levins’ own example helpfully illustrates the point. In the early population genetics models of Haldane, Fisher, and Wright, the environment was always assumed to be unchanging and static. This was regarded as a reasonable idealization and many biologists continue to use it to this day. Levins wants to criticize this particular idealization in contemporary work, but make room for the use that the early theorists made of it. He does this by arguing that for Haldane, Fisher, and Wright’s theoretical interests, treating the environment as static was an acceptable idealization.

Fisher, Wright, and Haldane’s interests involved characterizing the central properties of the major evolutionary forces. This involved many simplifications which were necessary in order to isolate these forces and study them individually. For example, they wanted to know how effective weak selection would be in bringing about evolutionary change. They did this by building models where natural selection was the only evolutionary force, a highly idealized assumption. Levins thinks that this was acceptable given the questions the founders of population genetics wanted to ask. Their work didn’t require getting the dynamic details of the environment correct. However, Levins thinks that many contemporary projects have goals that do require including these details. (Levins, 1966, 19)

Context sensitivity in model building is actually part of a larger theme that Levins and Richard Lewontin have emphasized. In the concluding chapter of *The Dialectical Biologist*, Levins and Lewontin claim that a commitment to what they call the dialectical world-view

involves understanding that the correct division of a whole phenomenon into parts depends on our theoretical interests. (Levins and Lewontin, 1985) What this means in the context of modeling is that how we break up a phenomenon into parts, and what features we choose to include and not to include in our model, must be a function of our modeling goals. There will not be a single norm that tells us how we should idealize in every case. Although this is a fairly extreme view and is probably not shared among many modelers, it seems to motivate many of Levins' discussions, especially in his response to Orzack and Sober 1993.

3.1. DESIDERATA

Levins' remarks about the historical grounding of theories are followed by his discussion of the desiderata of modeling and the different strategies associated with them. He claims that there are four *desiderata* of model building: manageability, generality, realism, and precision, but only the latter three are discussed. Although well-illustrated with examples from the literature, Levins' discussion does not contain definitions for the desiderata. So we have to rely on his examples and related discussion to determine the meanings of these terms.

Levins' notion of generality is best illustrated in the contrasts he draws between different kinds of models. One helpful contrast is the difference between Levins' own highly general models of evolution in changing environments and the fishery models of Watt 1956. Watt and co-workers constructed very complex mathematical models which include many factors known to have some effect on the phenomena of interest. They tailored their models to particular target phenomena by including specific causal factors associated with these systems and assigning weights to these causal factors from measured quantities. Because of their level of detail and specificity, the models only apply to a small number of phenomena.

Levins' models are quite different. Only a few causal factors are included in these models and many of the parameters are specified very imprecisely, leaving vague the magnitude of the causal forces described by the model. However, these models apply to many systems because few assumptions were made about the exact nature of the target system being modeled.

Working from this example, Levins' desideratum of generality is roughly the number of target systems that a model can be applied to. Yet this notion of generality is ambiguous between two senses of the term: how many actual target systems a model describes and how many logically possible systems a model describes. We can call these

*a-general*ity and *p-general*ity respectively. (Weisberg, 2003; Weisberg, 2004) Most of the time, Levins writes as if ‘generality’ means the number of actual target phenomena a particular model applies to. However, many of his examples seem to be general in the second sense. Consider one-locus models of the sort described by Wright and Fisher. One way to look at these models is that they are approximations of how many real phenomena behave. This makes them general in the actual-systems sense. But one-locus models also have a role in carving up logical space. They tell us something about possible, but non-actual, phenomena with a certain causal structure will behave.

Since there is no passage in the text that explicitly endorses one of these types of generality, I suggest we simply leave this as an open question. Generality can be understood either in the *p-general* or the *a-general* sense, and in fact, different uses of models will require different senses of generality.

3.1.1. *Realism*

Like generality, Levins’ use of the term ‘realism’ is also ambiguous. In some passages, he uses the term ‘realism’ as a synonym for accuracy. In others, it is related to diverse considerations including the the number of factors included in the model and the standards we use to evaluate the model.⁵

Even confining ourselves to defining realism in terms of accuracy, an ambiguity remains. Is realism an assessment of how well the structure of the model represents the structure of the world or is it an assessment of how close the *output* of the model matches some aspect of the target phenomenon? In Levins discussion of the second strategy of model building, it seems like he had in mind a match between the causal structure of the world in and the mathematical structure of the model. He talks about the structural features of simple Volterra predator-prey models as not being very good representations of the world because they omit time lags, physiological states, and density effects. ?19]rl:strategies However, in his discussion of the first strategy of model building, Levins conceives of realism as predictive accuracy. He discusses how fishery models can be used to make very accurate predictions and takes this to be a mark of realism. They are realistic even though the structural features of these models are not necessarily good representations of the world. Both senses of accuracy are important and there is really no need to choose one or the other until we find ourselves engaged in a detailed discussion of particular tradeoffs or the value of a particular kind of model.

⁵ Elaborate on differences between 1966 and 1993 versions.

3.1.2. *Precision*

Precision ought to be the least controversial of Levins' three desiderata. This is because precision is a well-understood statistical property, formally defined as "the closeness of repeated measurements of the same quantity." (Sokal and Rohlf, 1981, 13) This makes precision a property of data, not models or other representations. Levins, however, describes precision as a property of models or the equations used to describe models, not data. So we cannot simply import the definition of precision from statistics.

Levins use of the term 'precision' is actually closer to the ordinary language use of the term where precision means something like "finesse of specification." This sense of precision can be given a statistical interpretation by associating precision with the number of significant figures, or error involved in a measurement.

Although Levins talks about sacrificing or gaining precision in modeling, it is best to understand precision as a property of *model descriptions*, the equations or other representations which specify models. (Weisberg, 2003; Weisberg, 2004; Weisberg, forthcoming-b; Godfrey-Smith, ms; Godfrey-Smith, ming) Precision is thus the fineness of specification of the parameters, variables, and other parts of model descriptions.

4. Three Strategies of Model Building

Maximizing the three desiderata of precision, realism, and generality might be desirable, but, Levins argues, it is impossible. In what has become the most well known part of "The Strategy," Levins argues that when building models in population biology, one can only maximize two of the three desiderata simultaneously. Hence this section of the paper is often described as Levins' statement of the tradeoff between realism, precision, and generality.

Levins does not actually use the term 'tradeoff' anywhere in the paper. In fact, in a later article 1993, Levins suggests that he was only making an observation about the state of population biology in the late 1960s, not making a logical claim that the three desiderata cannot be maximized. Nevertheless, I think the article does presume some kind of three-way tradeoff between these desiderata. Reading the text carefully even provides us with some clues about the nature of the tradeoff Levins envisioned. He says that we cannot simultaneously *maximize* precision, realism, and generality. This does not mean that it is impossible to increase the magnitude of all three properties simultaneously. This is a *weak tradeoff*; it is not a zero-sum game situation where you can only

increase the magnitude of one of these properties by decreasing another. If you started from some very low magnitude of the three desiderata, you could increase the three of them simultaneously. However, if you start from a maximum degree of two of the three desiderata, it will be impossible to increase the third, without simultaneously decreasing the magnitude of one of the others.

Levins neither defends the existence of the three-way tradeoff nor does he even use the term. However, almost everyone who has cited “The Strategy” has taken this to be Levins’ intent. It is therefore natural to ask whether or not there really is a three-way tradeoff between realism, generality, and precision.

Interpreted literally, Levins’ three-way tradeoff does not exist. Ironically, this is because of the existence of a two-way tradeoff between precision and generality. Precision tradeoffs off against p-generality and it is impossible to *increase* the magnitude of these properties simultaneously. It is a *strong tradeoff*. Strictly speaking there is no tradeoff between precision and a-generality. Increasing precision makes the achievement of a-generality more difficult, but not impossible. (Weisberg, 2003) Since it is impossible to increase precision and p-generality simultaneously, then one possible situation envisioned by Levins’ three-way tradeoff — maximizing p-generality and precision — is impossible. Thus Levins’ three-way tradeoff cannot really exist.

Although I am skeptical about Levins’ three-way tradeoff interpreted literally, I do think that his identification of a “tradeoff” reveals important relationships obtaining between the properties of models. For example, there are several determinate relationships between precision and generality such as the one discussed above. Similarly, there are relationships between precision and accuracy, which is one part of what Levins’ calls realism. While preliminary studies have revealed some of the tradeoffs and other relationships between the properties of models, much is still unknown. A full analysis of these properties would be an important addition to our knowledge of the nature of theories and theory construction.

One of the most important reasons to develop an analysis of model properties and their relationships is that this analysis will help us develop an account of rational idealization. I think that this is the great insight of Levins’ discussion. To see how this program could be carried out, let’s provisionally accept the existence of Levins’ tradeoff which he uses to describe three research strategies which are meant to be alternatives to the brute-force approach. Each of these strategies can be thought of as embodying a different representational ideal. If we enumerate representational ideals on the basis of such constraints and analyze the the value of these different ideals, we can develop a theory

of rational idealization. Such an approach is just under the surface of Levins' article and this, I believe, is the enduring importance of what Levins accomplishes in "The Strategy."

Levins' three strategies for model building fall straightforwardly out of the three desiderata involved in the tradeoff. We can 1) sacrifice generality to gain precision and realism, 2) sacrifice realism to gain generality and precision, or 3) sacrifice precision to gain generality and realism.

4.1. FIRST STRATEGY: BUILD PRECISE AND REALISTIC MODELS

The first strategy involves building models that are very realistic and precise, but lack generality. These models are most useful to build when one's primary interest is in making quantitative predictions about the short-term behavior of a system. Levins cites fishery biologists (e.g. Watt, 1956) as being the major practitioners of this strategy.

Constructing such models requires detailed and precise knowledge about the system one is trying to model, for there are many parameters that need to be set. After making many precise measurements, one feeds these results into the description of the model in the form of parameters. When one wants to make a prediction about a particular system, one also specifies the initial conditions very precisely. Since realistic models of complex phenomena are mathematically unwieldy, these models will most likely require manipulation on a computer. Theorists feed their measurements into a computer and are able to generate highly precise predictions about very specific phenomena.

Levins has few general criticisms of this kind of modeling, perhaps because he thinks its strengths and weaknesses are self-evident. Such an approach generates highly accurate models useful for forecasting in limited domains. On the other hand, they tell us little about analogous systems and aren't very good at isolating the most important factors that affect the evolution of the phenomena they model.

The first strategy of modeling is extremely similar to the brute-force approach, so much so, in fact, that they may not be distinguishable. Levins may believe that this strategy isn't a form of the brute-force approach because it explicitly endorses the sacrifice of generality. However, as I understand it, brute-force modeling only requires complete representation of a particular phenomena; it has no requirement of generality. Thus I actually think that Levins' first strategy of modeling is a variant the brute-force approach and in his discussion of it, he

simply gives us more details about the problems associated with this approach.⁶

4.2. SECOND STRATEGY: BUILD GENERAL AND PRECISE MODELS

Sacrificing realism to gain generality and precision is the second strategy of model building. If you are willing to make a large number of approximations and work with a highly abstract representation, you can generate precisely specified models that will apply approximately to many target systems. This approach is especially conducive to analyzing a model's structure in significant mathematical detail.

An example of this strategy is the logistic model of population growth. If we take the following model description:

$$\frac{dN(t)}{dt} = (r - \frac{r}{K}N(t))N \quad (1)$$

and specify a precise value for r and K , we can describe the dynamics of this system exactly. Such a model is not very realistic in that there are many factors that affect population growth which are not included in the model. This means that even if the model makes accurate predictions about the growth of a population, it could only be thought to represent the actual process taking place if we set our standards of adequacy low, allowing highly approximate and abstract representations. But this is just what we expect when adopting the second strategy, the model is supposed to be highly general and very precisely specified, but not necessarily an accurate description of any particular phenomenon.

Levins associates the second strategy of modeling with the physicists who have entered population biology. Whether or not this continues to be true today, for many population biologists adopt something like the second strategy, Levins' deeper point is that these models are analogous to the idealized models used in physics. For one thing, these simple models are analogous to "frictionless planes" for biological systems. The comparison to the physical sciences is also apt in that modelers adopting the second strategy have been able to borrow mathematical structures directly from physics. For example, population biologists use a form of the diffusion equation to model genetic drift. This equation was first developed to deal with particles diffusing out in to space or

⁶ One possible way for Levins to avoid the conflation of the BFA and the first strategy is to point out that there are different loci for generality. One may be committed to a research program that is highly general, but which will require brute-force models for individual phenomena. So the techniques can remain general, while the individual models are not. Perhaps this allows for a brute force strategy that is not, strictly speaking, the first strategy of model building.

among other particles, but can also be used to model genetic drift and other stochastic biological processes. (Roughgarden, 1979, 69)

One of Levins' criticisms of the brute-force approach to modeling is that it generates models which do not admit of analytical solutions. In the more general version of the criticism that I presented, brute-force models often do not allow for detailed mathematical analyses. If one is especially concerned about this problem, then models generated using the second strategy have special appeal. Not many models in population biology admit of analytical solutions, but the ones that do are almost always generated using the second strategy, such as the one-locus model of natural selection, the logistic growth model, and the Volterra predator-prey model. Although only a small percentage of models generated using this strategy will admit of analytical solutions, it tends to generate models with a very definite mathematical structure on which mathematical analysis can be performed. This has enabled important breakthroughs in population biology in areas as diverse as fundamentals of evolutionary theory (e.g., Fisher, 1930) and the biomedical sciences (e.g., Nowak and May, 2000).

Before moving to the third strategy, we should ask whether one can really adopt the second strategy as characterized by Levins. I have already discussed the fact that there is a strong tradeoff between precision and generality. Doesn't this make the second strategy an impossible one to pursue?

Strictly speaking, according to the way I have defined precision on behalf of Levins, the second strategy would be quite difficult to pursue. One could find an intermediate balance between precision and generality and perhaps simultaneously increase the magnitude of both of these properties by attenuating realism somehow, but still, you couldn't pursue a strategy of gaining maximal generality and maximal precision. The best you could do would be to maximize the conjunction of generality and precision.

Yet there are clearly a cluster of models that Levins has identified with the second strategy. These models are very simple, their structures are often borrowed from the physical sciences, and they are often analyzed as algebraic structures alone, with no values being specified for their parameters.

If precision means fineness of specification of the parameter values, then not specifying parameters actually makes the description maximally imprecise. This would have the result, of course, of maximizing generality. So ironically, maybe the best way to characterize the second strategy of modeling is: sacrifice realism and be maximally imprecise about the values of the parameters (i.e. algebra only) in describing the model, so as to maximize generality. This correctly identifies the

physics-like models that Levins identifies with the second strategies and shows why it is rational to pursue it.

4.3. THIRD STRATEGY: BUILD REALISTIC AND GENERAL MODELS

The third strategy of model building involves sacrificing precision, to gain realism and generality. Such models, according to Levins, are highly flexible and are often described and analyzed graphically rather than algebraically. The parameters of the model descriptions associated with these models are often specified in qualitative, not quantitative, ways. In describing such a model, theorists might specify that a particular quantity is increasing or decreasing, but wouldn't specify the exact numerical value for the parameter which controlled the shape of the curve. They might specify that a polynomial term in the description is convex or concave, but not give the exact value to the coefficient of this term. The results of these models, at least in population biology, are sometimes expressed as inequalities, or using some other kind of limiting functional form.

The third strategy of modeling is clearly Levins' preferred mode of theorizing. It figures prominently in his best known work from the period (e.g. 1962, 1966) and he was interested in offering it as a viable alternative to the second strategy, which he associated with the highly unrealistic assumptions made by physicists studying biological systems. Yet surprisingly, Levins' doesn't defend the special status of the third strategy nor does he give any explanation of its value in contrast to the other strategies.

There are several reasons that one might adopt the third strategy. The simplest reason is that precision is the least intrinsically valuable desideratum. If one needs to make a highly precise prediction, then precision will be of value. But if not, there is no particular reason that precision matters in scientific inquiry. Generality and realism, however, are important for a much more fundamental goal of scientific inquiry, giving scientific explanations. Realism is required for explanations because we cannot explain something if we cannot even characterize it accurately. Generality is required, according to some accounts, so that similar, but distinct phenomena can be explained in a parallel way, leading to greater unification in our explanations. (Friedman, 1974; Kitcher, 1981) These are clearly advantages of the third strategy and may be why Levins seems to favor the strategy above the others.

4.4. TOWARD A THEORY OF IDEALIZATION

I opened this section by claiming that the most important part of Levins' paper is not the tradeoff or even the strategies of modeling

he outlines. Rather, the most important part of the paper is the way Levins' analysis gives us a template for developing a theory of rational idealization. In concluding this section, let us consider how this is so.

Levins' analysis of three strategies modelers can adopt was driven by his identification of a three-way tradeoff he believed that they faced. A more expansive analysis of the strategies of modeling could take into account other properties of models and analyze the way in which these properties trade off or constrain one another. The analysis would begin by looking at properties similar to accuracy and generality, but could also go on to consider theoretical desiderata of specific special sciences or even of human cognitive capacities.

The next step would be to analyze the value and importance of each desideratum. Not all theoretical desiderata are intrinsically valuable and even those that are may be more important in particular scientific contexts than in others. Levins performs this kind of analysis at points in "The Strategies," but it is a project that needs to be carried out in much greater detail.

Once we understand the intrinsic and instrumental value of the various desiderata associated with modeling, we can combine this information with our knowledge of tradeoffs to build a theory about rational idealization. For example, if we believe, as Levins seems to, that generality is associated with explanatory power, we might recommend the construction of highly general models for explanations. We would consider the various ways to increase generality and develop strategies, such as Levins' third strategy, which would achieve this goal with low cost for other explanatorily important values.

5. Robustness

The final section of "The Strategy" discusses the notion of *robustness*, or what Wimsatt calls *robustness analysis*. Pre-theoretic intuition and traditional philosophy of science suggest that as theoretical practice matures, we will see a gradual unification and reduction of the number of models for particular target systems. However, if we take Levins' picture of the idealization approach seriously, we are assured that theoretical practice will continue to generate a proliferation of models. Since different models score better with respect to the different desiderata, we will expect scientists to continue generating multiple models for a given phenomenon. In addition, because all models will contain idealizations, within particular strategies, multiple models may also be generated incorporating different idealizations.

One might lament this fact and work toward the eventual elimination of multiple models in favor of a single, or small set of models, for a particular phenomenon, but Levins thinks that there is a special value in having multiple models. Since modelers deal in idealizations, all of their models are false. Yet when a model makes a prediction or purports to explain something, we need a way of determining “whether a result depends on the essentials of the model or on the details of the simplifying assumptions.” (1966, 20) Levins proposes that we do this by searching for *robust theorems*. “[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)”

The terminology is a little misleading here. A robust theorem is not a theorem as conventionally understood. It is a conditional statement linking a common structure of a set of models, to some behavioral or static property predicted by those models. Levins claims that when we discover a property or dynamic that is implied by multiple models, we can be more confident that this theorem does not depend on the idealizations we have made. Instead, it depends on core causal properties shared by the set of models.

Robustness analysis seems to be some kind of confirmation procedure, one that is applicable in situations where all of the models are false and highly idealized. For example, Levins considers the following statement to be a robust theorem:

In an uncertain environment species will evolve broad niches and tend toward polymorphism. (Levins, 1966, 20)

This result, he tells us, can be derived from three kinds of models: the fitness set model (Levins, 1962), a model using the calculus of variation, and a genetic model (Levins and MacArthur, 1966).⁷

Several attempts have been made to further articulate Levins notion of robustness. In a critical article by Orzack and Sober 1993 robustness analysis is taken to be a simple, non-empirical method of testing hypotheses. They envision the modeler looking at an exhaustive set of models and determining whether a particular hypothesis is a logical consequence of each of the models. Only in the special case where a theorem is a logical consequence of each and every model, they argue, is the hypothesis confirmed by robustness analysis.

It is unlikely that such an analysis could be carried out, for no theorist can really generate an exhaustive list of all possible models for a particular phenomenon. Even if such a list could be given, it

⁷ Subsequent research (e.g., Seger and Brockmann, 1987) has called in to question whether this particular phenomenon is actually robust.

is extremely unlikely that the very same theorem would be a logical consequence of each model. Orzack and Sober are clearly aware of this and hence are quite skeptical of robustness analysis. There is every reason to be skeptical of robustness analysis as characterized by Orzack and Sober, yet I believe that they have given an overly stringent interpretation of the procedure. Levins was not offering an alternative to empirical confirmation; rather, he was explaining a procedure used in conjunction with empirical confirmation in situations where one is relying on highly idealized models.

Another analysis of robustness was offered by Wimsatt 1981 who considers the technique to be part of a much broader set of procedures where one attempts to isolate the core components of a process. He writes that

[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. (Wimsatt, 1981, 128)

Thus, for Wimsatt, everything from “[u]sing different experimental procedures to verify the same empirical relationships” to “[u]sing different assumptions, models, or axiomatizations to derive the same result or theorem” to “[u]sing failures of invariance . . . to calibrate or recalibrate our measuring apparatus” are part of robustness analysis. (Wimsatt, 1981, 127)

I have a much less expansive understanding of Levins’ notion of robustness than Wimsatt does. As I read Levins, finding a robust theorem means finding that a set of diverse models imply the existence of the same dynamic or property. When we find such a robust property, robustness analysis involves examining the core structure of the models that implied this result. If there is some core causal structure in all of these models that gives rise to the robust behavior, we have confirmed, *ceteris paribus*, the *connection* between a certain causal structure and the robust property.

By showing that a dynamic or property is robust, we have shown that it is characteristic of the core causal structure. Whenever the causal structure which is common to the models is instantiated in a real target system, we can expect to see the robust property or dynamic. Finding that the robust theorem is implied by each member of a set of diverse models is necessary to show that this property or dynamic is

independent of the extraneous details of the models and only depends on the shared, core causal structure.

An example of this kind of robustness is the discovery of the Volterra Principle. The Volterra Principle says that “the result of a general insecticide will be an increase in the abundance of the pests, themselves, whereas the number of the predators on the pests will decrease.” (Roughgarden, 1979, 439) This principle was not discovered empirically, but by robustness analysis. Although it was first derived using the simple Volterra model of predator-prey interactions, it arises in many other models. Specifically, it “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)

If we understand the Volterra Principle as a robust theorem, then we can give a Levinsonian explanation of the principle’s significance as follows: A number of related, but distinct models of predator-prey interaction all describe a certain dynamic between predator and prey which ecologists now call the Volterra principle. This theorem is robust and comes about in any model that represents a certain causal structure. What this tells us is that in any possible system with this causal structure, regardless of other details, will be an instantiation of the Volterra principle. In this sense, the Volterra principle is confirmed via its being a robust theorem.

Robustness analysis is a procedure followed by many modelers and is thought to confer a confirmation-like status on robust hypotheses. This is something that has been neglected by most philosophers of science and we have a lot to learn by thinking carefully about it. While Levins deserves much credit for emphasizing the importance of the procedure, his original discussion gives us few clues as to how and why the procedure works. This, I suspect, has led critics like Orzack and Sober to be highly critical of the idea.

Differing from the analyses of Orzack and Sober, Wimsatt, and Levins (in response to Orzack and Sober), I think the robustness analysis is simply a way of isolating a core structure which is common to multiple models and then determining the consequences of this core structure. The procedure itself does not confirm hypotheses; rather, it identifies an already well-confirmed sub-structure contained in a number of models and shows what must follow from this consequence. The result is the partial confirmation of conditional statements such as “whenever you have positive coupling of kind P between predator and prey, you will expect a dynamic like D .” This is defended and explained in much greater detail in Weisberg, forthcoming-a, ms-a.

6. Conclusion

Few papers, scientific or philosophical, seem as relevant forty years after publication as they did when they were first published. Levins' insights about the nature of theoretical practice were more nuanced than many well-known philosophical accounts of theorizing from the 1960s. Perhaps this accounts for why it has taken the greater part of forty years to start incorporating Levins' ideas into accounts of modeling, confirmation, and idealization.

Undoubtedly Levins' article will always be remembered in connection with the three-way tradeoff between realism, generality, and precision. This is certainly an important reason to remember the paper and has led to a number of insightful discussions in both the biological and philosophical literature. However, as I hope to have made clear, the paper should also be remembered for its emphasis on the *practice* of theorizing, not just the outcome of theorizing. This takes us outside the comfortable realm of analyzing theories, and into the murkier realm of the goals and intentions of modelers.

The paper should also be remembered as an important lesson about how to construct an account of idealization. The articulation of the strategies of modeling and then an analysis of the advantages and disadvantages of these strategies is a crucial step in such an account. Finally, the article should be remembered for drawing our attention to the special issues that arise in the context of confirming highly idealized theories. Accounts of confirmation may need to make room for robustness analysis and the other techniques used for dealing with the proliferation of models for complex systems.

"The Strategy" and Levins work on theoretical methods is finally getting the attention it deserves by those who write about theories and theorizing. My hope is that when the students of my own students write "Eighty Years of 'The Strategy,'" philosophy of science will look quite different and will have fully absorbed what Levins began teaching us about theoretical methods forty years ago.

References

- Fisher, R. A.: 1930, *The Genetical Theory of Natural Selection*. Oxford: The Clarendon Press.
- Friedman, M.: 1974, 'Explanation and Scientific Understanding'. *Journal of Philosophy* **71**, 5–19.
- Godfrey-Smith, P.: forthcoming, 'The Strategy of Model Based Science'. *Biology and Philosophy*.
- Godfrey-Smith, P.: ms., 'Model-based Science and the Representational Theory of Mind'.

- Kitcher, P.: 1981, 'Explanatory Unification'. *Philosophy of Science* **48**, 507–531.
- Levins, R.: 1962, 'Theory of Fitness in a Heterogeneous Environment, I. The Fitness Set and Adaptive Function'. *American Naturalist* **96**(861), 361–373.
- Levins, R.: 1966, 'The Strategy of Model Building in Population Biology'. In: E. Sober (ed.): *Conceptual Issues in Evolutionary Biology*. Cambridge, MA: MIT Press, first edition, pp. 18–27.
- Levins, R.: 1968, *Evolution in Changing Environments*. Princeton: Princeton University Press.
- Levins, R.: 1993, 'A Response to Orzack and Sober: Formal Analysis and the Fluidity of Science'. *Quarterly Review of Biology* **68**(4), 547–555.
- Levins, R. and R. Lewontin: 1985, *The Dialectical Biologist*. Cambridge, MA: Harvard University Press.
- Levins, R. and R. H. MacArthur: 1966, 'The Maintenance of Genetic Polymorphism in a Spatially Heterogeneous Environment'. *American Naturalist* **100**, 585–589.
- Morgan, M. S. and M. Morrison: 1999, 'Models as Mediating Instruments'. In: M. Morgan and M. Morrison (eds.): *Models as Mediators*. Cambridge: Cambridge University Press, pp. 10–37.
- Nowak, M. A. and R. M. May: 2000, *Virus Dynamics*. Oxford: Oxford University Press.
- Odenbaugh, J.: ms., 'Complex Systems, Trade-Offs and Mathematical Modeling: Richard Levins' "Strategy of Model Building in Population Biology" Revisited'.
- Odum, H. T.: 1983, *Systems Ecology: An Introduction*. New York: Wiley.
- 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 Publishing Co.
- Seeger, J. and H. J. Brockmann: 1987, 'What is bet-hedging?'. *Oxford Surveys in Evolutionary Biology* **4**, 181–211.
- Sober, E. (ed.): 1984, *Conceptual Issues in Evolutionary Biology*. Cambridge, MA: MIT Press, first edition.
- Sokal, R. R. and F. J. Rohlf: 1981, *Biometry*. San Francisco: W.H. Freeman, second edition.
- Strevens, M.: 2004, 'The Causal and Unification Approaches to Explanation Unified — Causally'. *Noûs* **38**, 154–179.
- Suppe, F.: 1977, 'The Search for Philosophic Understanding of Scientific Theories'. In: F. Suppe (ed.): *The Structure of Scientific Theories*. Chicago: University of Illinois Press, 2nd edition.
- Suppes, P.: 1960a, 'A Comparison of the Meaning and Use of Models in Mathematics and the Empirical Sciences'. *Synthese* **12**, 287–300.
- Suppes, P.: 1960b, *Models of Data*, pp. 251–61. Stanford: Stanford University Press.
- Volterra, V.: 1926, 'Fluctuations in the Abundance of a Species Considered Mathematically'. *Nature* **118**(558–60).
- Watt, K. E. F.: 1956, 'The Choice and Solution of Mathematical Models for Predicting and Maximizing the Yield of a Fishery'. *Journal of the Fisheries Research Board of Canada* **13**, 613–345.
- Weisberg, M.: 2003, 'When Less is More: Tradeoffs and Idealization in Model Building'. Ph.D. thesis, Stanford University.
- Weisberg, M.: 2004, 'Qualitative Theory and Chemical Explanation'. *Philosophy of Science* **71**, 1071–1081.
- Weisberg, M.: forthcoming-a, 'Robustness Analysis'. *Philosophy of Science*.

- Weisberg, M.: manuscript-a, 'Three Kinds of Idealization'.
- Weisberg, M.: forthcoming-b, 'Who is a Modeler?'. *British Journal for Philosophy of Science*.
- Weisberg, M. and K. Reisman: manuscript-b, 'The Robust Volterra Principle'.
- Weisstein, E. W.: 2003, *CRC Concise Encyclopedia of Mathematics*. Boca Raton: Chapman & Hall, second edition.
- Wiegert, R. G.: 1975, 'Simulation Models of Ecosystems'. *Annual Review of Ecology and Systematics* **6**, 311–338.
- 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, pp. 124–163.
- Wimsatt, W. C.: 1987, 'False Models as a Means to Truer Theories'. In: M. Nitecki and A. Hoffmann (eds.): *Neutral Models in Biology*. Oxford: Oxford University Press, pp. 23–55.

