



University of Pennsylvania
ScholarlyCommons

Departmental Papers (Philosophy)

Department of Philosophy

9-2009

The Structure of Tradeoffs in Model Building

John Matthewson
Australia National University

Michael Weisberg
University of Pennsylvania, weisberg@phil.upenn.edu

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



Part of the [Philosophy Commons](#)

Recommended Citation

Matthewson, John and Weisberg, Michael, "The Structure of Tradeoffs in Model Building" (2009).
Departmental Papers (Philosophy). 16.
https://repository.upenn.edu/philosophy_papers/16

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/philosophy_papers/16
For more information, please contact repository@pobox.upenn.edu.

The Structure of Tradeoffs in Model Building

Abstract

Despite their best efforts, scientists may be unable to construct models that simultaneously exemplify every theoretical virtue. One explanation for this is the existence of tradeoffs: relationships of attenuation that constrain the extent to which models can have such desirable qualities. In this paper, we characterize three types of tradeoffs theorists may confront. These characterizations are then used to examine the relationships between parameter precision and two types of generality. We show that several of these relationships exhibit tradeoffs and discuss what consequences those tradeoffs have for theoretical practice.

Keywords

Tradeoff, model, modeling, biological models, richard levins, generality, precision, theoretical virtues

Disciplines

Philosophy

The Structure of Tradeoffs in Model Building*

John Matthewson & Michael Weisberg[†]

Draft of September 12, 2007(13f)
under review

Abstract

Despite their best efforts, scientists may be unable to construct models that simultaneously exemplify every theoretical virtue. One explanation for this is the existence of tradeoffs: relationships of attenuation that constrain the extent to which models can have such desirable qualities. In this paper, we characterize three types of tradeoffs theorists may confront. These characterizations are then used to examine the relationships between parameter precision and several types of generality. We show that several of these relationships exhibit tradeoffs and discuss what consequences those tradeoffs have for theoretical practice, especially in sciences that study complex phenomena.

1 Introduction

Much of the theoretical research in evolutionary biology, ecology, and other sciences studying complex systems proceeds via the study of mathematical models. For any given phenomenon of interest, one often finds myriad models differing in scope, generality, precision, accuracy, and the number

*Many thanks to Brett Calcott, Patrick Forber, Alan Hájek, Peter Godfrey-Smith, Kim Sterelny, Deena Skolnick Weisberg and especially Aidan Lyon, Ryan Muldoon, Daniel Singer, and Scott Weinstein for extraordinarily helpful comments. This research was partially supported by National Science Foundation grant SES-0620887 to MW.

[†]The authors contributed equally to this paper.

of included causal factors in the literature. Theorists often motivate this proliferation of models by pointing to the tradeoffs they face; they claim, for example, that they cannot maximize precision and generality, so some models are more precise and some more general.

Richard Levins argued that such tradeoffs require theorists to adopt different strategies of model building depending on their theoretical goals. (1966) His own discussion argued that a three-way tradeoff exists between generality, realism, and precision, such that a model builder cannot simultaneously maximize all of these desiderata. Recent philosophical literature has raised doubts regarding Levins' three-way tradeoff (Weisberg, 2006) and has even questioned the existence of tradeoffs in general (Orzack & Sober, 1993; Orzack, 2005). However, the bulk of the new literature on this topic argues that at least some tradeoffs are real, be they psychological, pragmatic (Odenbaugh, 2003), or the result of the logic of representation (Weisberg, 2004).

In this paper, we reexamine the concept of tradeoffs discussed by Levins, by biologists working in Levins' tradition, and by philosophers of science. We argue that there is not one, but at least three relevant types of tradeoff. After giving definitions for these, we investigate their interrelationships and consider the circumstances under which one type of tradeoff implies another. Finally, the paper revisits the relationship between generality and precision, two desiderata that have previously been alleged to exhibit a tradeoff. With our new taxonomy of tradeoffs, we show that precision and generality do in fact trade off, and that the relationship between these modeling attributes is more restrictive than was previously argued.

2 Tradeoffs

Tradeoffs are relationships of *attenuation* that hold between two or more modeling attributes, or what Levins (1966) called *desiderata* of model building. Attenuation occurs when an increase in the magnitude of one attribute makes the achievement of another more difficult. In this paper, we will confine ourselves to discussions of two-way attenuation. Attenuation comes in

at least four varieties, only three of which are actually tradeoffs and hence of greatest concern to us in this article.

The first kind of attenuation is *simple attenuation*. Two attributes exhibit simple attenuation if and only if increasing the magnitude of one modeling attribute makes it more difficult but not impossible to increase the magnitude of another attribute. Simple attenuation imposes important pragmatic constraints on modelers, as the resources required to deal with the attenuation may be considerable. However, there are also forms of attenuation that cannot be overcome through further data collection or other resources. Specifically, we identify and define three such types: *strict tradeoffs*, *increase tradeoffs*, and *Levins tradeoffs*.

2.1 Strict Tradeoffs

We begin with a procedural definition for strict tradeoffs. Two desiderata exhibit a strict tradeoff (s-tradeoff) when an increase in the magnitude of one desideratum necessarily results in a decrease in the magnitude of the second, and vice versa. In other words, as the magnitude of one of the desiderata goes up, the other must come down.

In order to formalize this concept, we will define the tradeoffs in terms of the set of possible values the modeling attributes can take when the tradeoff applies. To do this, we begin by defining a set Λ , which contains ordered tuples corresponding to the possible magnitudes that different attributes can take in a single model. Because we will only be considering pairwise relationships between modeling attributes, we will designate Λ to be the set of possible simultaneous values for two modeling attributes P and Q . Thus $\Lambda = \{\langle p_i, q_i \rangle\}$, where each ordered pair $\langle p_i, q_i \rangle$ is a pair of allowable simultaneous magnitudes for P and Q .

When no tradeoff obtains between P and Q , then Λ will correspond to the cartesian product of the possible magnitudes of P and Q . But when a tradeoff obtains, only specific pairs will be included in Λ . In the case of a strict tradeoff, the only ordered pairs allowed in Λ are those that satisfy the following constraint: for every set of ordered pairs, if P in the first pair is

smaller than P in the second pair, then Q in the first pair must be larger than Q in the second pair.

Since we will need to refer to elements of Λ , we adopt the notation that π and ρ are arbitrary elements of Λ , each designating a different *pair* of values that can be associated with a particular model. $(\pi)_1$ designates the first element in ordered pair π , and $(\pi)_2$ designates the second. For example, if we are considering the magnitudes for precision and generality that are associated with a model, π denotes a particular possible precision/generality pair for that model, and $(\pi)_1$ denotes the precision of the model as expressed by the pair and $(\pi)_2$ the generality. The symbols $<$ and $>$ indicate an ordering over these elements in the standard way. If we wish to compare the precision designated by two possible pairs where the value given for precision in one is greater than in the other, this may therefore be symbolized $(\pi)_1 < (\rho)_1$. Using this notation, we can give the following definition for a strict tradeoff.

Definition 1. *Let $\Lambda = \{\langle p_i, q_i \rangle\}$, where each $\langle p_i, q_i \rangle$ corresponds to a pair of possible simultaneous magnitudes for P and Q . Let π and ρ be two distinct elements of Λ . A **strict tradeoff** obtains between P and Q iff $\forall \pi \forall \rho [(\pi)_1 < (\rho)_1 \leftrightarrow (\pi)_2 > (\rho)_2]$*

This definition can be applied to a wide range of modeling attributes, because the attributes in question need not share any scale or continuity properties. It is only required that the two attributes can be expressed as ordered pairs of simultaneously achievable magnitudes. While the definition generalizes beyond easily graphable cases, it is useful to represent the tradeoff graphically as we have in Figure 1.

Although the definition is stated in terms of $(\pi)_1$ having smaller magnitude than $(\rho)_1$, it is fully symmetric because the quantifiers on π and ρ both range over each ordered pair. Thus, the definition entails that when the magnitude of one of these two attributes goes down, the magnitude of the other must go up. This symmetry is important because, if it were allowable to decrease one magnitude without a corresponding increase in the other, any reversal of this alteration would result in an increase in the magnitude




Figure 1: Diagram of a strict tradeoff representing the possible magnitudes allowed for attributes P and Q as a dashed line. π and ρ are two arbitrary points, which are elements of Λ . Along with all of the other allowable points, they satisfy the definition of a strict tradeoff. The negative slope of this line is a signature of strict tradeoffs, although the definition also applies to cases where a slope is not well-defined.

of that attribute without an associated decrease in the other. This is of course precisely what is prohibited according to our informal definition of the s-tradeoff.

The definition also stipulates that Λ contains at least two distinct elements to avoid cases where the biconditional is trivially satisfied. These cases are excluded because tradeoffs are only of scientific interest when they occur due to the interaction of modeling attributes.

2.2 Increase Tradeoffs

The second kind of tradeoff is an increase tradeoff, or i-tradeoff. Informally, two modeling attributes exhibit an i-tradeoff when the magnitude of the

attributes cannot both be increased simultaneously. That is, an increase in the magnitude of one cannot be accompanied by an increase in the other. In order to formalise this definition we will once again invoke a set of ordered pairs Λ , defined as above.

Definition 2. Let $\Lambda = \{\langle p_i, q_i \rangle\}$, where each $\langle p_i, q_i \rangle$ corresponds to a pair of possible simultaneous magnitudes for P and Q . Let π and ρ be two distinct elements of Λ . An **increase tradeoff** obtains between P and Q iff $\forall \pi \forall \rho [((\pi)_1 < (\rho)_1 \rightarrow (\pi)_2 \geq (\rho)_2) \& ((\pi)_2 < (\rho)_2 \rightarrow (\pi)_1 \geq (\rho)_1)]$

The difference between the i-tradeoff and the s-tradeoff is that in the i-tradeoff case there can be subsets of Λ where the magnitude of one of the modeling attributes increases as we move between elements, but the magnitude of the second stays the same. This represents a significant constraint on a modeler, but one less stringent than the s-tradeoff. This definition also stipulates that one cannot decrease both attributes in the presence of an i-tradeoff. Figure 2 illustrates a situation that exhibits an i-tradeoff but not an s-tradeoff.

2.3 Levins Tradeoffs

Two attributes exhibit a Levins tradeoff (L-tradeoff) when the magnitude of both of these attributes cannot be simultaneously maximized. If we look in Λ for the maximum value of P (p_{\max}) and the maximum value of Q (q_{\max}), then when an L-tradeoff obtains, there is no ordered pair $\langle p_{\max}, q_{\max} \rangle$ in Λ .

Definition 3. Let $\Lambda = \{\langle p_i, q_i \rangle\}$, where each $\langle p_i, q_i \rangle$ corresponds to a pair of possible simultaneous magnitudes for P and Q . Let π be an element of Λ . Further, let p_{\max} be the maximum value for P in Λ and q_{\max} be the maximum value for Q in Λ . A **Levins tradeoff** obtains between P and Q iff $\neg \exists \pi [((\pi)_1 = p_{\max}) \& ((\pi)_2 = q_{\max})]$

We call this a ‘Levins tradeoff’ because it is closest to the concept of a tradeoff discussed in Levins’ philosophical work (Weisberg, 2006). Levins tradeoffs are only of significance when both model attributes in question




Figure 2: Diagram of an increase tradeoff representing the possible magnitudes allowed for attributes P and Q (i.e. the points in Λ) as a dashed line. π and ρ are arbitrary points in Λ that satisfy the definition of an i-tradeoff. The graphical signature of increase tradeoffs is a non-positive slope, although the definition applies in cases where slope is not well-defined.




Figure 3: Illustration of a Levins tradeoff between two attributes. π lies on the horizon of the tradeoff, illustrated by the dashed line. In this particular case, there is a region of positive value about the simultaneous in-principle maxima that the model attributes cannot reach, designated by the shaded area. Apart from this limitation, the attributes of the model can take any combination of magnitudes.

have a maximum magnitude in Λ , otherwise a Levins tradeoff vacuously obtains.

Whenever an L-tradeoff holds between two attributes, there exists a set of possible combined values that the attributes cannot exceed, which we call the *horizon*. We use the term ‘horizon’ to differentiate this in-practice limit from the in-principle maxima the attributes could simultaneously attain if the L-tradeoff did not hold. Considered graphically, the horizon defines the upper limit above which the magnitudes of the two attributes cannot *simultaneously* extend (see Figure 3).

The inability to simultaneously maximize two modeling desiderata means that if a modeler wishes to maximize the magnitude of one desideratum, she must accept that the magnitude of the other will be suboptimal. The upshot of this is that when faced with an L-tradeoff, the modeler must make strategic decisions regarding which attribute to maximize. A model that is constructed under such constraints will therefore be determined at least in part by the goals of the modeler in question.

Similarly, and perhaps more acutely, the existence of s-tradeoffs and i-tradeoffs will strongly influence decisions regarding what kind of model will be best for the job at hand, possibly eliminating the utility of some models outright. For example, if an s-tradeoff makes maximization of either attribute too costly in terms of the other, it may be that for the model to be of any use, *neither* attribute can be maximized. In Section 5, we argue that an s-tradeoff obtains between one kind of generality and precision. Since precision and generality are often both important in scientific practice, it will often be the case that any useful model will have intermediate magnitudes for both of these desiderata.

3 Relationships Between the Tradeoffs

We have now given formal definitions for the three kinds of tradeoffs. The interrelationships between the tradeoffs are relatively complex, but there are some clear entailments among them. In this section, we will argue that the existence of an s-tradeoff entails that an i-tradeoff and an L-tradeoff occur,

and that when a weak condition is met, the existence of an i-tradeoff entails that an L-tradeoff occurs.

3.1 Strict Tradeoffs

The existence of an s-tradeoff between two attributes entails an i-tradeoff between those attributes. This follows trivially from the formal definitions of i- and s-tradeoffs. Since \leq is equivalent to ($<$ or $=$), $A < B$ entails that $A \leq B$. As the definition for an increase tradeoff is identical to the definition for a strict tradeoff with \leq instead of $<$ in the consequent of each conditional, any instance that satisfies the criteria for a strict tradeoff will also satisfy the criteria for an increase tradeoff.

The case can be made intuitively as follows: The existence of a strict tradeoff means that as we increase the magnitude of A , we must decrease the magnitude of B . If we must decrease the magnitude of one attribute whenever we increase the magnitude of the other, we cannot increase the magnitude of both attributes. Therefore an increase tradeoff will also hold.

When an s-tradeoff holds between two attributes, then an L-tradeoff will also hold between those attributes. The idea is intuitive, but the justification of this entailment is more complex, so we give a formal proof below. Note that the entailment only goes through due to the specification that an s-tradeoff requires two distinct elements, otherwise if Λ contained only a single element $\langle p_{\max}, q_{\max} \rangle$, this would trivially exhibit an s-tradeoff but would not exhibit an L-tradeoff.

Theorem 1. *The existence of an s-tradeoff between two attributes entails an L-tradeoff between those attributes*

Proof. Let $\Lambda = \{\langle p_i, q_i \rangle\}$, where each $\langle p_i, q_i \rangle$ corresponds to a pair of possible simultaneous magnitudes for P and Q . Let p_{\max} be the maximum value for P in Λ and q_{\max} be the maximum value for Q in Λ .

Assume that P and Q exhibit an s-tradeoff. This means that $\forall \pi \forall \rho [(\pi)_1 < (\rho)_1 \leftrightarrow (\pi)_2 > (\rho)_2]$ Assume that an L-tradeoff does not obtain. This means that $\langle p_{\max}, q_{\max} \rangle$ is an element of Λ . Let X designate this element. Therefore $(X)_1 = p_{\max}$ and $(X)_2 = q_{\max}$.

For a strict tradeoff to hold, there must be at least two distinct elements in Λ . Let Y designate an arbitrary element of Λ that is distinct from X . We can instantiate the universal quantifiers in the definition of an s-tradeoff using the elements X and Y , first instantiating X for π and Y for ρ , then Y for π and X for ρ , to give us the formulae: $[(X)_1 < (Y)_1 \leftrightarrow (X)_2 > (Y)_2]$ and $[(Y)_1 < (X)_1 \leftrightarrow (Y)_2 > (X)_2]$

Subtheorem 1. $(X)_1 = (Y)_1$

Proof. $[(Y)_1 < (X)_1 \leftrightarrow (Y)_2 > (X)_2]$ can be satisfied if either both $(Y)_1 < (X)_1$ and $(Y)_2 > (X)_2$ are true, or if they are both false. It is clear that they cannot be both true, as we know that $(X)_2 = q_{\max}$, and therefore it is impossible for $(Y)_2 > (X)_2$ to be true.

For both to be false, it must be the case that $(Y)_1 < (X)_1$ is false, so either $(Y)_1 > (X)_1$ or $(Y)_1 = (X)_1$ must be true.

Since $(X)_1 = p_{\max}$ it must be that $(X)_1 = (Y)_1$ □

Subtheorem 2. $(X)_2 = (Y)_2$

Proof. If we follow the same reasoning as for the case above, but applied to the $[(X)_1 < (Y)_1 \leftrightarrow (X)_2 > (Y)_2]$ instantiation, then we see it must be that $(X)_2 = (Y)_2$ □

X and Y are distinct elements of Λ , and therefore cannot have exactly the same values for both members in the ordered pair. But we have proved that for an s-tradeoff to occur in the absence of an L-tradeoff, it must be that both $(X)_1 = (Y)_1$ and $(X)_2 = (Y)_2$, which results in a contradiction. Thus, if an s-tradeoff holds between two attributes, this entails that an L-tradeoff also holds between those attributes. □

3.2 Increase Tradeoffs

We have seen that whenever an s-tradeoff holds between two attributes, this entails that an i-tradeoff and an L-tradeoff also hold between those at-

tributes. Since i-tradeoffs are strictly weaker than s-tradeoffs, the existence of an i-tradeoff does not entail an s-tradeoff. However, it can be shown that the existence of an i-tradeoff entails an L-tradeoff as long as there is at least one member in Λ where neither attribute is maximal. If we know that an i-tradeoff holds between two attributes and this weak condition is met, no member in Λ can have both attributes at maximum, and therefore an L-tradeoff will hold between those attributes.

Theorem 2. *The existence of an i-tradeoff and any element with submaximal values for both attributes entails an L-tradeoff*

Proof. Let $\Lambda = \{\langle p_i, q_i \rangle\}$, where each $\langle p_i, q_i \rangle$ corresponds to a pair of possible simultaneous magnitudes for P and Q . Let p_{\max} be the maximum value for P in Λ and q_{\max} be the maximum value for Q in Λ . Assume that P and Q exhibit an i-tradeoff. Applying our definition, we know that $\forall \pi \forall \rho [((\pi)_1 < (\rho)_1 \rightarrow (\pi)_2 \geq (\rho)_2) \& ((\pi)_2 < (\rho)_2 \rightarrow (\pi)_1 \geq (\rho)_1)]$.

Assume that one element in Λ has sub-maximal values for both P and Q . Let Y designate this element. In that case, $(Y)_1 < p_{\max}$ and $(Y)_2 < q_{\max}$.

Now assume that P and Q do not exhibit an L-tradeoff. This means that $\langle p_{\max}, q_{\max} \rangle$ is also an element of Λ . Let X designate this element. In that case, $(X)_1 = p_{\max}$ and $(X)_2 = q_{\max}$. We can instantiate our definition of an i-tradeoff with $[(Y)_1 < (X)_1 \rightarrow (Y)_2 \geq (X)_2] \& [(Y)_2 < (X)_2 \rightarrow (Y)_1 \geq (X)_1]$

Because both attributes are submaximal in Y , $(Y)_1 < (X)_1$. According to our definition, this means that $(Y)_2 \geq (X)_2$, which is impossible since $(X)_2 = p_{\max}$ and $(Y)_2$ is submaximal. This is a contradiction.

Therefore the existence of an i-tradeoff between two attributes plus the existence of at least one element in Λ that has submaximal values for both of these attributes entails the existence of an L-tradeoff between those attributes.

□

3.3 Levins Tradeoffs

In some cases, the existence of an L-tradeoff and the restriction of Λ to certain sets of values may entail other tradeoffs. For example, if the magnitudes of two attributes of a given model lie on the horizon of an L-tradeoff, this will lead to a “partial” increase tradeoff, as it is impossible to increase both magnitudes from any point on the horizon. However, we do not believe that any other form of attenuation is entailed by the existence of an L-tradeoff without such further limitations.

Regardless of this, the existence of an L-tradeoff in and of itself can shape and constrain scientific practice in significant ways. In particular, the existence of an L-tradeoff means that it is impossible for a single model to maximize every desirable attribute. It was in this spirit that Levins argued from the existence of tradeoffs to a particular strategy of model building. He claimed that if no single model can be optimal in all respects, theorists ought to construct multiple models, each of which sacrifices one desideratum in order to maximize others. These individually restricted models can then be compared and integrated in an attempt to meet our competing modeling goals.

While we doubt that Levins tradeoffs arise for logical and representational reasons in the absence of s- and i-tradeoffs, we believe that scientists confront them on a regular basis. In such cases, the L-tradeoffs are domain-specific and depend on empirical facts, not logic.

Particularly striking examples of L-tradeoffs can be found in the uncertainty principles of both classical and quantum mechanical wave mechanics. In quantum mechanics, for example, there is an L-tradeoff between the specification of a particle’s position and its momentum. At a fine enough grain of resolution, it becomes impossible to maximize the precision of both of these quantities. In fact, the L-tradeoff can be specified precisely as $\Delta x \Delta p_x \geq \frac{1}{2} \hbar$. This constraint means that any model which requires the specification of position and momentum, such as the ones developed to account for chemical phenomena like reaction rates, cannot maximize the precision of specification for both of these quantities (Levine, 1991).

This L-tradeoff between the specification of position and momentum has nothing to do with the logic of representation, but rather the fact that position and momentum are described by particular hermitian operators and that these operators do not commute. The fact that position and momentum are accurately described by these operators in quantum mechanics was an empirical discovery. Had the world been different and these quantities were described by commuting operators, then they would be simultaneously knowable with maximal accuracy. This is thus an example of a logically contingent L-tradeoff.

3.4 Summary of the tradeoffs and their interactions

A strict tradeoff obtains between two modeling attributes when it is impossible to increase the magnitude of one without a decrease in the magnitude of the other, and it is impossible to decrease one without an increase in the other. We have shown that if an s-tradeoff holds between two attributes, this entails that an i-tradeoff and L-tradeoff also hold between those attributes. An i-tradeoff obtains between two attributes when it is impossible to increase the magnitudes of both attributes and it is impossible to decrease the magnitudes of both attributes. If an i-tradeoff holds between two attributes, and at least one pair of these attributes has neither at maximum value, then an L-tradeoff also holds between those attributes. An L-tradeoff holds between two attributes when it is impossible to simultaneously maximize both attributes.

4 Precision and Generality

Having looked in detail at the three kinds of tradeoffs and some of the relationships between them, we now analyze the relationships that obtain between precision and generality. We have chosen to focus on these desiderata for this paper because of their role in earlier discussions of tradeoffs and because of the connections they have to some of the larger goals of theoretical practice, such as wide descriptive breadth, the discovery of similarities

across disparate systems, and increased explanatory power. We will return to these connections in further detail at the end of the paper. In this section, we focus only on the tradeoffs between these desiderata.

Earlier analyses have argued for the existence of tradeoffs between precision and different forms of generality (Levins, 1966, 1968; Weisberg, 2004, 2006). We agree that there are tradeoffs between these desiderata, but will show that the attenuation is more stringent than was previously believed. To do this, we begin by clarifying how the relevant terms are defined.

4.1 Precision

When discussed in the context of mathematical modeling, *precision* is often characterized as *parameter precision*, an attribute of the equations that describe mathematical models. Unlike the statistical notion of precision, which is an attribute of data, parameter precision is a measure of the fineness of specification of parameter values.

Defining precision for an entire mathematical model description requires an assessment of the precision of each of the individual parameters in that description. We can do this by first defining the *uncertainty* associated with some parameter value as the deviation of that value from the best estimate of the true value.

Canonically, a parameter value might be written as the central value for the parameter plus or minus the uncertainty associated with it. The central value is often interpreted as a best guess or best estimate, so we can write the canonical value of an imprecisely defined parameter value as $p_{best} \pm \delta p$, where δp is the uncertainty.

Precision can then be defined in terms of uncertainty as follows:

Definition 4. *If a parameter p has value $p_{best} \pm \delta p$, then that parameter's **precision** is $1/2\delta p$.*

Precision is defined as the reciprocal of two times the uncertainty to preserve the intuitive idea that precision increases as uncertainty decreases.

It becomes a more complex issue to define precision when dealing with multiple parameters. Probably the best general way to aggregate parameter

precision is with the use of an n -dimensional distance formula. Because these details are not necessary for the discussion in this article, we will rely on the following comparative test of precision applicable to the cases discussed in this paper:

A model description D_1 is more *precise* than a model description D_2 if D_1 picks out a proper subset of the models picked out by D_2 , when all other factors are held fixed.

This test relies on the fact that when we compare two equations with different degrees of precision, the more precise one will describe a subset of the models described by the less precise one. This method only works when comparing equations that have the same number of parameters and obey the nesting relations discussed below. Most measures of parameter precision will be unable to judge which description is more precise if these conditions do not hold.

To illustrate these ideas, we can consider the family of models of exponential population growth. An *uninstantiated* model description, or equation, that picks out this family can be written as follows

$$\frac{dN}{dt} = rN \tag{1}$$

In this equation the variable N stands for population size, and r stands for the growth rate of the population, which is assumed to be constant in the models designated by this description.

This equation can be instantiated with differently specified values for r , picking out different sets of models. For example, r might be instantiated to give the descriptions:

$$\frac{dN}{dt} = (1.5 \pm 0.1)N \tag{2}$$

$$\frac{dN}{dt} = (1.50 \pm 0.01)N \tag{3}$$

Changes in precision will effect the set of models picked out by a descrip-

tion. More precise descriptions will pick out subsets of the sets of models picked out by less precise descriptions. This is easiest to show in cases such as (2) and (3) above, where the parameter values in the different descriptions overlap. Description (2) picks out all of the models that have a value for r between 1.4 and 1.6, while description (3) only picks out the models with a value for r between 1.49 and 1.51. Note that all of the models picked out by description (3) are also picked out by description (2), but not vice versa. This means that the set of models picked out by description (3) forms a proper subset of the models picked out by description (2).

4.2 Fidelity and Model-world relations

The next modeling attribute we will discuss is generality. At first pass, this may seem straightforward: Generality is a measure of how many phenomena a model or set of models successfully relate to. However, the manner in which models relate to their targets is not a simple or uncontroversial issue. Before we proceed, we need to be clear regarding how we view the model-world relationship in this context.

One position regarding this relationship is that, ideally, a relationship of isomorphism should hold between a model and a mathematical representation of its target (van Fraassen, 1980). Another possibility is that the relationship is less stringent but still model-theoretic, such as homomorphism (Lloyd, 1994) or partial isomorphism (da Costa & French, 2003). Other possibilities include metric relationships, such as distance in a state space between model and target, as well as less formal notions of similarity, such as the ones offered by Cartwright (1983), Giere (1988), and Godfrey-Smith (2006). Pluralism about these matters is, of course, also possible (Downes, 1992).

While our own view on these matters is probably closest to the pluralism advocated by Downes, we believe that what follows is compatible with all of these ideas, except perhaps the most stringent reading of the demand for isomorphism. Rather than taking sides in this debate, we will use the term ‘applies to’ when describing the relationship between model and target

phenomenon. We do this to emphasize something that is not a typical focus in the literature. What we want to highlight are the standards a modeler brings to bear when determining whether the model applies to a target. In other words, not only is it important to assess the fidelity of a particular model, which might be assessed with model-theoretic, metric, or informal similarity measures, it is also important to understand the standards of fidelity applied by the modeler. We call these standards *fidelity criteria*.

There are many ways in which scientists may employ models, and a modeler's fidelity criteria will vary according to the task at hand. Any particular model may be of use to different investigators in different ways and for different reasons, according to which aspects of the target and model concern them. For example, two theorists may have different demands for how predictively accurate a model is: one might require near-exact results, while another might be satisfied if the model only approximately predicts the behavior of the target. Alternatively, another theorist might be solely concerned with modeling the causal dependencies within the target, and have very limited requirements regarding fidelity in terms of predictive accuracy.

From this, we see that modelers may be concerned with the outcome of the model: its predictions about the quantities of measurable attributes. We call this the *dynamical fidelity* of the model and the theorists' standards for dynamical fidelity are the *dynamical fidelity criteria*. Alternatively or additionally, a theorist might be interested in how well a model describes the causal structure of the target system. The assessment of this attribute is the *representational fidelity* of the model and the criteria for assessing it are the *representational fidelity criteria*. (Weisberg, 2007)

For example, meteorologists require high dynamical fidelity from their models, because they need to make precise and accurate predictions of the weather. However, because the goal is to predict rather than explain, meteorologists may not require high representational fidelity from these models. Conversely, theorists constructing macroevolutionary models often aim for high representational fidelity, even if they cannot predict observables with quantitative accuracy. Models with biomedical applications may require high degrees of both kinds of fidelity so that effective medical interventions

can be planned. Usually, a model will require a complex mixture of these two types of fidelity.

The fidelity criteria in use in any modeling situation will have a notable effect on the generality of the models under consideration. All else being equal, more permissive fidelity criteria will tend to mean that a given model will apply to more phenomena. For this reason, in our analyses of the interaction between precision and generality we will assume that the fidelity criteria used to assess the relationship between models and target are held fixed. In section 5.3, we will revisit how changing fidelity criteria can affect the tradeoffs.

4.3 Generality

Unlike precision, which is an attribute of the equations we use to describe models, the concept of generality we are interested in concerns the model-target relationship. In what follows, we take generality to be a measure of how many targets the models in question apply to.¹ Having made this clarification, there are two further ways in which the concept of generality must be disambiguated.

First we need to differentiate between generality regarding the number of target systems an individual model applies to, and regarding the number of target systems a set or family of models applies to. In the first instance, if model m_1 applies to more targets than model m_2 , then m_1 is more general than m_2 . In the second, if the set of models M_1 as a whole applies to more targets than the set of models M_2 , then set M_1 is more general than set M_2 .

These two types of generality can co-vary; increasing the generality of the individual models in a given set will often also increase the generality of the set as a whole. However, individual model generality and model set generality can also come apart. For example, we will show that it is possible to increase model set generality while holding individual model generality fixed. The fact that individual and set generality can come apart means that

¹Note that we do not use ‘generality’ to refer to the inclusiveness of causal factors in the model. As we discuss in section 5.3, we use the term ‘scope’ to refer to this property.

they must be analyzed separately when we consider whether they tradeoff against other modeling attributes.

Generality must also be disambiguated according to the type of target systems considered. It must be clear whether we are concerned with how many *actual* targets a model or set of models applies to, or how many *possible* targets the model or set applies to.² We call these different types of generality *a-generality* and *p-generality* respectively. P-generality is not something that only philosophers might take seriously; it is often what scientists have in mind when they discuss how general a model is, especially in the context of its explanatory power. Sometimes exploration of the non-actual helps explain the actual, and the point of some explanatory models is not necessarily to resemble any real systems, but to canvass possibility space. For example, biological models that generalize to show that the fitness costs for a species to have three sexes are too high can explain why there are no such species in the real world (Fisher, 1930).

P- and a-generality will usually take different values for any given model or set of models; the number of actual targets applied to will be different from the number of possible targets applied to except in some extremely limited circumstances. Additionally, two models might be identically p-general but differ in their a-generality and vice versa. Once again, this means that these different types of generality will have different relations to other modeling attributes, and will therefore need to be considered separately in an analysis of tradeoffs.

As we sometimes measure generality in terms of logical possibility, we will be dealing with infinite sets of possible targets. This means that we cannot always order the generality of models or sets of models according to cardinality, but will have to consider whether they apply to some set of target systems that is a proper subset of another, thus being of lesser generality. This is a less universal measure than we might like, since it

²We will interpret ‘possible targets’ to mean *logically possible* targets. One might also use the term to pick out *nomologically* or *physically* possible targets. We prefer the broader, logical interpretation because the interests of modelers range from what is known to be actual through what is known to be physically impossible. Future analyses of tradeoffs might fruitfully explore more restricted modalities.

restricts us to cases where we are comparing sets that stand in set/subset relations to each other, but to date this is most comprehensive way we know of to analyze p-generality.

Since individual model and model set generality can take the p- and a-form, we have four types of generality and hence four interactions to analyze. Each will be considered in turn in the following section.

5 Tradeoffs Between Precision & Generality

We now have the tools in place to assess the relationship between precision and generality in the context of modeling. As stated above, we begin our analysis by isolating precision and generality, holding everything including the fidelity criteria fixed. After this analysis, we will consider what occurs when the fidelity criteria are allowed to vary, arguing that this can affect the generality of a model or set of models, which can in turn modify what results from the tradeoff.

5.1 Precision and P-generality

First we consider how an increase in precision affects individual model p-generality. Recall that precision is an attribute of model descriptions, not of models themselves. Alterations in precision modify the number of models picked out by a description, but not how these models apply to their targets. As all other features such as fidelity criteria are held fixed, this means that the number of logically possible targets any given individual model will apply to is unchanged. Individual model p-generality is therefore unaffected when precision is manipulated.

However, as noted previously, a more precise model description picks out a proper subset of the models picked out by a less precise counterpart. Since each individual model applies to a fixed number of logically possible targets, the set of logically possible targets applied to by a subset of models is also a subset compared to the targets applied to in the less precise case. This means that model set p-generality is decreased whenever precision is

increased.

We can consider this argument in detail:

1. Assume model description d picks out a set of models M_1 .
2. If model description d' is more precise than d , d' will pick out M_2 , a set of models that is a proper subset of M_1 .
3. Since all attributes other than precision are held fixed, each individual model applies to the same number of possible target systems as previously.
4. This means that, since M_2 is a proper subset of M_1 , the models in M_2 apply to a proper subset of the logically possible target systems applied to by M_1 .
5. Therefore, by definition, M_2 is less p-general than M_1 .
6. Therefore, increasing the precision of a model description means that model set p-generality is reduced.

The preceding argument shows that any increase in precision will impose a cost on p-generality. However, recall that two attributes only exhibit a strict tradeoff when an increase in either attribute results in a decrease in the other. Therefore in order to assess whether precision and p-generality exhibit a strict tradeoff, we must check to see if the attenuation is symmetrical. We can do this with the reverse argument:

1. Assume model description d picks out a set of models M_1 and model description d' picks out a set of models M_2
2. If M_2 is more p-general than M_1 , M_2 must apply to a superset of the possible targets M_1 applies to.
3. Since all attributes other than precision are held fixed, each individual model applies to the same number of possible target systems as previously.

4. This means that, since the models in M_2 apply to a superset of the logically possible target systems in M_1 , it must also be the case that M_2 is a superset of M_1 .
5. Therefore d' is less precise than d .
6. Therefore, increasing the p-generality of a set of models means that the precision of the model description is reduced.

If all other attributes are held fixed, the only way we can increase p-generality is to decrease precision. Since we have shown that an increase in either precision or generality imposes a cost on the other, the relation between precision and model set p-generality is a strict tradeoff.

5.2 Precision and A-generality

Next we turn to the relationship between precision and a-generality. This is more complex than the p-generality case, since the effect that an alteration in precision has on a-generality will at least in part be determined by the empirical features of the particular system under consideration. We have seen that an increase in precision entails a decrease in p-generality. However, the actual targets that a model applies to will nearly always be far fewer than the logically possible targets the model applies to, so a reduction in p-generality does not necessarily imply a reduction in a-generality. This means that an increase in precision will only sometimes come at the expense of a-generality, dependent on the systems modeled and the attributes of those targets that are of interest to the modeler.

This is particularly clear when we consider the difference between how changes in precision affect the a-generality of models used in disciplines whose typical targets are homogenous with respect to the properties of interest, and those whose targets are heterogeneous. In both the homogeneous and heterogeneous cases, increases in precision may or may not lead to the exclusion of any actual targets. However, there will be a limit to how precise a model description can be before any actual targets are necessarily

excluded. The more homogeneous the target systems of interest, the more precise the description can be before this limit is reached.

For example, the targets and attributes that ecological models are directed towards are often very heterogeneous. The intrinsic growth rate, the attribute corresponding to r in our population growth model description, can be extremely varied from population to population. Consider the difference between the growth rate in a population of cane toads (*Bufo marinus*) in Australia and in the Americas. In Australia, cane toads have multiplied so rapidly as to constitute an ecological disaster, while in their native habitats they are largely static in numbers. Cases such as these mean that a model description that contains a finely specified value for r will often pick out a set of models that only applies to a small proportion of the relevant target populations. Precisely specified values of r will correspond to models appropriate only for studying the dynamics of either Australian cane toads or American cane toads, but not cane toads in general.

On the other hand, models in the physical sciences are often directed towards homogeneous sets of targets. This allows for the possibility of writing down highly precise model descriptions, which pick out a relatively small set of models, and yet still apply to all of the target systems of interest. For example, a model description describing models of electron mass and charge can be extremely precise, yet still pick out a set of models that applies to all electrons. This possibility results from the fact that physical quantities such as the mass and charge of fundamental particles are extremely homogeneous. In terms of these attributes, at least, there are no differences between the electrons in any part of the world, or for that matter the electrons on Alpha Centuri and those on Earth.

Because the degree of homogeneity of the target systems alters the effect that an increase in precision has on the a-generality of sets of models, the exact relationship between these attributes will vary on a case-by-case basis and requires specific empirical information about the targets being modeled. That said, there are some general features we can point to regarding the interaction between precision and a-generality.

Consider the relationship between precision and individual model a-

generality. As in the the p-generality case, changes in precision have no effect on individual model a-generality. Precision only determines whether a given model is picked out by a given description, not how the models relate to targets. So as long as we hold all other attributes fixed, any given model that is picked out will apply to the same number of actual targets regardless of changes in precision. Therefore changes in precision have no effect on individual model a-generality.

A second general relationship concerns precision and model set a-generality. Regardless of the system modeled, it is impossible to increase both precision and a-generality if all other attributes are fixed. We know that an increase in precision means that the set of models picked out by a description applies to a subset of the logically possible targets compared to previously. As discussed above, whether this will reduce the number of actual targets applied to depends on the systems themselves; however, we know a priori that there is no way that the targets in the world could be arranged such that reducing the size of our set of models while keeping all else fixed could increase the number of actual targets to which our set of models applies. This means that we cannot increase both precision and a-generality.

We can show this with a similar argument to that used above:

1. Assume model description d picks out a set of models M_1 .
2. If model description d' is more precise than d , d' will pick out M_2 , a set of models that is a proper subset of M_1 .
3. Since all other attributes are held fixed, each individual model applies to the same number of actual targets as previously, even if this is zero.
4. Therefore, as M_2 is a proper subset of M_1 , M_2 cannot apply to more actual targets than M_1 .
5. Therefore by definition, M_2 cannot be more a-general than M_1 .
6. Therefore, it is not possible to increase precision and also increase model set a-generality.

Again, in order to show that a tradeoff obtains between precision and a-generality, we need to check that the relationship holds in both directions. However, this time a simple reversal of the previous argument does not give us a symmetrical outcome, as the attenuation from a-generality to precision is stronger than from precision to a-generality.

1. Assume model description d picks out a set of models M_1 and description d' picks out M_2 .
2. If M_2 is more a-general than M_1 , M_2 must apply to more actual targets than M_1 .
3. Since all other attributes are held fixed, each individual model applies to the same number of actual targets as previously, even if this is zero.
4. This means that the only way that M_2 can apply to more actual targets than M_1 is if M_2 is a superset of M_1 .
5. If M_2 is a superset of M_1 , then d' is less precise than d .
6. Therefore, increasing model set a-generality means that precision must be decreased.

When we increase the precision of a model description, we cannot simultaneously increase a-generality, and if we increase a-generality, we must decrease precision. This means that precision and a-generality do not exhibit a strict tradeoff, as a necessary cost is incurred in only one direction. However, because an increase in a-generality incurs a cost in precision, this makes it impossible to increase a-generality and also increase precision. This means that a simultaneous increase is impossible in both directions, and so a-generality and precision exhibit an i-tradeoff.

We have now assessed the tradeoff relations between precision and the four categories of generality. Precision and both types of individual model generality show no tradeoffs. Precision and model set p-generality exhibit a strict tradeoff, and precision and model set a-generality exhibit an increase tradeoff.

5.3 The Role of Scope and Fidelity Criteria

Precision and model set a-generality always exhibit an i-tradeoff, but the case of the cane toads given above illustrates that when a high degree of heterogeneity among targets is present, the attenuation is often stronger. Here we will show that in addition to the intrinsic degree of heterogeneity between targets, changes in the evaluative standards adopted by the modeler can also strengthen or weaken the tradeoff between precision and a-generality. Very roughly, sets of models will apply to fewer cases when one evaluates them with higher standards of fidelity or intends that they capture more aspects of the target phenomena. These increased demands concern the fidelity criteria and *intended scope* respectively.

Recall that fidelity criteria are the standards used to assess the degree of similarity between model and target. These criteria are set by individuals and communities of modelers. They may have very relaxed fidelity criteria regarding a certain aspect of model and target, interested only in qualitative aspects of that particular property. Alternatively, they may have demanding fidelity criteria, and require that the model gives outputs that mimic the target very closely. The choice of these criteria typically depends on how the model will be used.

The *intended scope* of a model refers to which aspects of the target the model is intended to capture. As noted in §4.3, we do not use ‘scope’ synonymously with generality. For us, the term refers not to the breadth of applicability of the model, but rather to the aspects of a given target or targets the theorist wants to capture with the model.

A modeler may be interested in modeling only a few aspects of her target and disregard the others. In this case, her intended scope is very restricted. Alternatively, a modeler may be inclusive regarding which properties of the particular target she wishes to model, in which case the model will have a broad intended scope.

For example, consider the Lotka-Volterra model of predator-prey dynamics, which only takes predator and prey growth and death rates, rate of prey capture by predators, and efficiency of predators at converting captures to

new predator births into account. Typically, a biologist who employs this model in her research intends the model to have limited scope. The model is not intended to capture the effects of population structure, alternative sources of food for the predators, the possibility of prey seeking cover, etc.³ However, she might widen the intended scope to include such factors, which would require addition of the appropriate parameters. It is important to note, therefore, that all things being equal, increasing the scope of a model will tend to make one's assessment of the model-target relationship more demanding: There are more ways in which the model will be tested against its target.

Returning now to the tradeoff between precision and a-generality, we can ask how changes in scope and fidelity might affect this tradeoff in particular cases. We begin by considering changes in intended scope. Our original contention was that when a set of target systems is highly heterogeneous, all things being equal, the attenuation between precision and a-generality will often be more costly. This effect can be weakened, however, if the modeler chooses a restricted intended scope. By choosing a limited scope, the modeler reduces the number of aspects of targets her models must apply to. In so doing, she can limit the *effective heterogeneity* of the targets, limiting the effects of increasing precision on a-generality.⁴

Consider the case of an ecologist modeling foraging behavior in a rainforest. If she attempts to capture the foraging behavior of all the species in a particular region of a rainforest, it will be impossible to use precise descriptions (which thereby pick out a limited number of models) without imposing heavy costs in a-generality. This is because there is a great deal of variation between the strategies of different species, making the properties of interest very heterogeneous. However, if she is selective regarding which properties of these complex targets she wants to model, disparate targets in

³Their presence might have an affect on the dynamical fidelity of the model, but this is a separate issue.

⁴This is not an absolute rule, of course, because it depends on both the kinds of heterogeneity in the targets and the particular scope restrictions employed by the modeler. In many cases, however, the result of limiting scope is to make the targets effectively more homogeneous.

the ecosystem can look more similar and therefore can be captured with a smaller set of models. For example, her scope might be restricted to the energetic aspects of the foraging. Since these factors rely on biochemistry and the distribution of resources in the ecosystem, individual differences among organisms will be considerably diminished.

Another dimension along which the effective heterogeneity of targets can be adjusted is via the fidelity criteria. When these criteria are lowered, small differences between targets become less relevant and, depending on the degree of heterogeneity among the targets, can be made negligible. This can result in each individual model becoming more a-general, and by extension, sets of these models becoming more a-general. In this way, the effects of the tradeoff between precision and model set a-generality can be modified due to the interaction of fidelity criteria and generality.

In conclusion, precision and model set a-generality always exhibit an i-tradeoff, but the disparity between the intended targets, combined with considerations of scope and fidelity can make the attenuation relationship between them more or less costly.

6 Tradeoffs in Scientific Modeling

Theorists face many methodological constraints, some of which will abate with improvements in technology or greater available resources. However, there are at least three constraints that will not dissipate with scientific progress, specifically, the tradeoffs discussed in this paper. Modeling desiderata can exhibit strict or increase tradeoffs due to facts regarding logic and representation alone, while Levins tradeoffs result from contingent empirical facts in particular domains.

As an example of these tradeoffs, we have shown that there is an s-tradeoff between precision and p-generality, and an i-tradeoff between precision and a-generality. Because of the entailments among the tradeoffs, this also means that precision and p-generality exhibit an increase tradeoff and a Levins tradeoff. In addition, by virtue of the fact that precision and a-generality exhibit an i-tradeoff, we know that whenever any model

description that is not maximally precise picks out a set of models that is not maximally a-general, *no* description and set of models in this setting can be maximally precise and maximally a-general.

We believe these results are significant for a number of reasons. First, a rigorous demonstration of these tradeoffs answers an earlier demand in the literature by Orzack and Sober (Orzack & Sober, 1993; Orzack, 2005) for demonstrable proof of the existence of tradeoffs. We have provided such a proof for tradeoffs among non-trivial modeling desiderata. We also believe that the definitions and techniques described in this paper can be used to demonstrate the existence of a number of other tradeoffs among properties such as simplicity, representational capacity, and accuracy, as well as precision and generality.

The particular case of tradeoffs between precision and generality we discussed are also important for scientific methodology, as they are connected to some of the interests and goals employed by modelers studying complex systems. In particular, generality has been alleged to be a key ingredient in scientific explanations and to help capture trends among similar, but distinct phenomena.

Generality features prominently in many philosophical theories of scientific explanation, and in some accounts, generality is considered to be the core of a scientific theory's explanatory power. Unificationists correlate explanatory power of a theory with the number of phenomena that it can subsume (i.e., its a-generality), modulo other constraints (Kitcher, 1981, 1989; Friedman, 1974). More recent causal accounts (e.g., Strevens, 2004; Woodward, 2003) rely on a notion akin to p-generality to find the optimal causal explanation of a particular event.

The value of generality for explanation has also been defended by scientists, notably population biologists. These scientists point to the especially high value of general models in allowing theorists to explain similar phenomena using the same framework in order to reveal underlying patterns among these phenomena (e.g., Roughgarden, 1979, 1997; May, 2001; Nowak & May, 2000). Increasing the generality of a set of models, perhaps by lowering precision, lets theorists treat these systems in a common framework.

In so doing, theorists may have a greater ability to determine the underlying features common to these systems, features which may be responsible for understanding patterns of interest.

These accounts all suggest that increases in generality are, *ceteris paribus*, associated with an increase in explanatory power. The existence of tradeoffs between precision and generality indicates that one way to increase an explanatorily valuable desiderata is by sacrificing precision. Conversely, increasing precision may lead to a decrease in explanatory power via its effect on generality.

Whether a particular instance of this tradeoff does in fact limit explanatory power will, of course, depend on the circumstances under which precision is altered. Arbitrarily altering generality by way of altering precision will not thereby necessarily increase or decrease explanatory power, as this depends on other factors as well. Similarly, at a certain point, a model can become so general that it says nothing of interest regarding its targets and thereby loses explanatory power. However, there are circumstances where it would be rational for theorists to sacrifice some precision to gain generality and hence explanatory power.

Therefore, if an increase in either explanatory power or descriptive breadth are of importance to the modeler, whenever they are faced with a heterogeneous set of targets they must make a choice. One way to achieve the desired levels of generality is to make the equations used to describe the model less precise than would otherwise be optimal. Alternatively, the way in which the model is compared with its targets may be altered, by limiting the intended scope or fidelity criteria employed by the modeler.

More broadly, this discussion gives us further reason to follow Levins in seeing the analysis of tradeoffs as crucial to understanding scientific methodology. An appreciation of what kinds of tradeoffs can occur and the circumstances in which they arise will aid philosophers in understanding the patterns of models used in the different branches of science.

References

- Cartwright, N. (1983). *How the laws of physics lie*. Oxford: Oxford University Press.
- da Costa, N. C. A., & French, S. (2003). *Science and partial truth*. Oxford: Oxford University Press.
- Downes, S. M. (1992). The importance of models in theorizing: A deflationary semantic view. *PSA 1992, 1*, 142–153.
- 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.
- Giere, R. N. (1988). *Explaining science: A cognitive approach*. Chicago: University of Chicago Press.
- Godfrey-Smith, P. (2006). The strategy of model based science. *Biology and Philosophy, 21*, 725–740.
- Kitcher, P. (1981). Explanatory unification. *Philosophy of Science, 48*, 507–531.
- Kitcher, P. (1989). Explanatory unification and the causal structure of the world. In *Minnesota studies in the philosophy of science, vol. 13*. Minneapolis: University of Minnesota Press.
- Levine, I. N. (1991). *Quantum chemistry* (fourth ed.). Englewood Cliffs, NJ: Prentice Hall.
- Levins, R. (1966). The strategy of model building in population biology. In E. Sober (Ed.), *Conceptual issues in evolutionary biology* (First ed., pp. 18–27). Cambridge, MA: MIT Press.
- Levins, R. (1968). *Evolution in changing environments*. Princeton: Princeton University Press.
- Lloyd, E. A. (1994). *The structure and confirmation of evolutionary theory* (second ed.). Princeton: Princeton University Press.
- May, R. M. (2001). *Stability and complexity in model ecosystems*. Princeton: Princeton University Press.
- Nowak, M. A., & May, R. M. (2000). *Virus dynamics*. Oxford: Oxford

University Press.

- Odenbaugh, J. (2003). Complex systems, trade-offs and mathematical modeling: Richard Levins' "Strategy of Model Building in Population Biology" revisited. *Philosophy of Science*, 70, 1496–1507.
- Orzack, S. H. (2005). What, if anything, is 'The Strategy of Model Building in Population Biology'? A comment on Levins (1966) and Odenbaugh (2003). *Philosophy of Science*, 72, 479–485.
- Orzack, S. H., & Sober, E. (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.
- Roughgarden, J. (1997). *Primer of ecological theory*. Upper Saddle River, NJ: Prentice Hall.
- Strevens, M. (2004). The causal and unification approaches to explanation unified — causally. *Noûs*, 38, 154–179.
- van Fraassen, B. C. (1980). *The scientific image*. Oxford: Oxford University Press.
- Weisberg, M. (2004). Qualitative theory and chemical explanation. *Philosophy of Science*, 71, 1071–1081.
- Weisberg, M. (2006). Forty years of "The Strategy": Levins on model building and idealization. *Biology and Philosophy*, 21(5), 623–645.
- Weisberg, M. (2007). Who is a modeler? *British Journal for Philosophy of Science*, 58, 207–233.
- Woodward, J. (2003). *Making things happen*. Oxford: Oxford University Press.