# The Structure of Tradeoffs in Biological Model Building*

John Matthewson & Michael Weisberg

*under review*

**Abstract**

This paper examines the tradeoffs alleged to hold between modeling attributes that are faced by biological modelers. Three distinct kinds of tradeoffs are defined and analyzed: strict tradeoffs, increase tradeoffs, and Levins tradeoffs. Employing these definitions and analyses, the putative tradeoff between precision and generality is reexamined and found to be more restrictive than previously reported. Parameter precision exhibits a strict tradeoff with p-generality and an increase tradeoff with a-generality.

## 1 Introduction

Much of the theoretical literature in evolutionary biology and ecology examines the different attributes of mathematical models. For any given phenomenon of interest, one finds a myriad of models differing in scope, generality, precision, accuracy, and the number of included causal factors. Theorists often motivate this proliferation of models by pointing to the tradeoffs they face; one cannot maximize precision and generality, they argue, 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

*Many thanks to Brett Calcott, Patrick Forber, Alan Hájek, 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.

precision, such that a theorist cannot simultaneously maximize all these desiderata. Recent philosophical literature has raised doubts regarding Levins' three-way tradeoff (Weisberg, 2006) and even questioned the existence of tradeoffs in general (Orzack & Sober, 1993). 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 reviews the tradeoff between generality and precision, which has been the most thoroughly explored in the literature to date. With our new taxonomy of tradeoffs, we show that the relationship between these modeling attributes is more restrictive than was previously thought.

## 2 Tradeoffs

Tradeoffs are relationships of *attenuation* that hold between two or more modeling attributes, or what Levins 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 three types of symmetrical attenuation that we call modeling tradeoffs: *strict tradeoffs*, *increase tradeoffs*, and *Levins tradeoffs*. We begin by defining each of these.

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 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 $\Lambda$, which contains ordered tuples corresponding to the possible magnitudes that the modeling attributes can take simultaneously. Because we will only be considering pairwise relationships between modeling attributes, we will designate $\Lambda$ 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 $\Lambda$ 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 $\Lambda$. In the case of a strict tradeoff, ordered pairs are only allowable if they satisfy the constraint that for any two pairs, if the magnitude for an attribute in the first pair is smaller than its magnitude in the second pair, then the magnitude of the other attribute in the first pair must be larger than its magnitude in the second pair. Adopting the notation that $\pi$ and $\rho$ are arbitrary elements of $\Lambda$, and $(\pi)_1$ the first element in ordered pair $\pi$, and that the symbols $<$ and $>$ denote an ordering over these elements in the standard way, 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 $\pi$ and $\rho$ be two distinct elements of $\Lambda$. 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 even continuity properties. The definition only requires that they are able to be ordered into 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 $(\rho)_1$, the definition is fully symmetric. It entails that when the magnitude of either attribute goes down, the magnitude of the other must go up. This further symme-

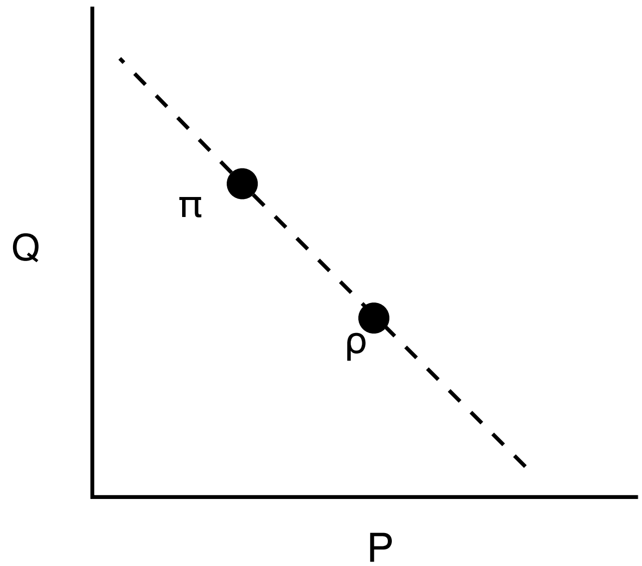


Figure 1: This diagram represents the possible magnitudes allowed for attributes $P$ and $Q$ as a dashed line. Each pair of points satisfies the definition of a strict tradeoff, such that for $\pi$ and $\rho$ if $(\pi)_1 < (\rho)_1$, then $(\pi)_2 > (\rho)_2$ and vice versa. The negative slope of this line is a signature of strict tradeoffs.

try 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 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 $\Lambda$ 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 because of the interaction of modeling attributes.

The second kind of tradeoff is an increase tradeoff, or i-tradeoff. 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 procedural definition we will once again invoke a set of ordered pairs $\Lambda$, defined as above.

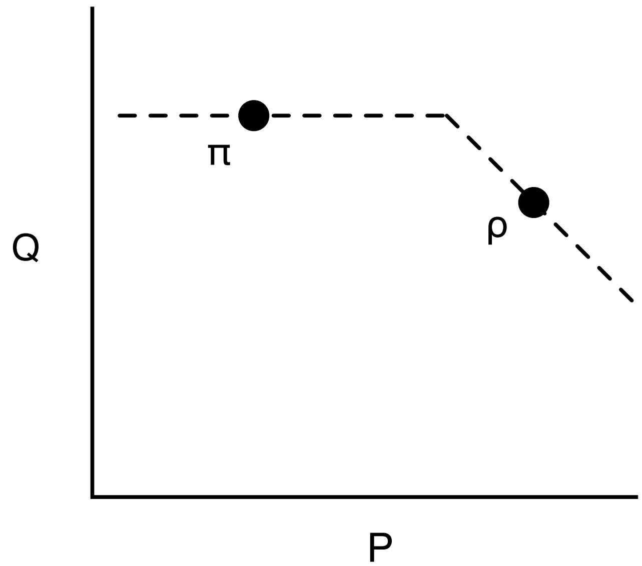


Figure 2: This diagram represents the possible magnitudes allowed for attributes $P$ and $Q$ as a dashed line. Each pair of points satisfies the definition of an increase tradeoff, such that for $\pi$ and $\rho$ if $((\pi)_1 < (\rho)_1 \rightarrow (\pi)_2 \geq (\rho)_2)$ and $((\pi)_2 < (\rho)_2 \rightarrow (\pi)_1 \geq (\rho)_1)$. The graphical signature of increase tradeoffs is non-positive slope.

**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 $\pi$ and $\rho$ be two distinct elements of $\Lambda$. 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 $\Lambda$ where the magnitude of one of the modeling attributes increases as we move between elements, but the second magnitude 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.

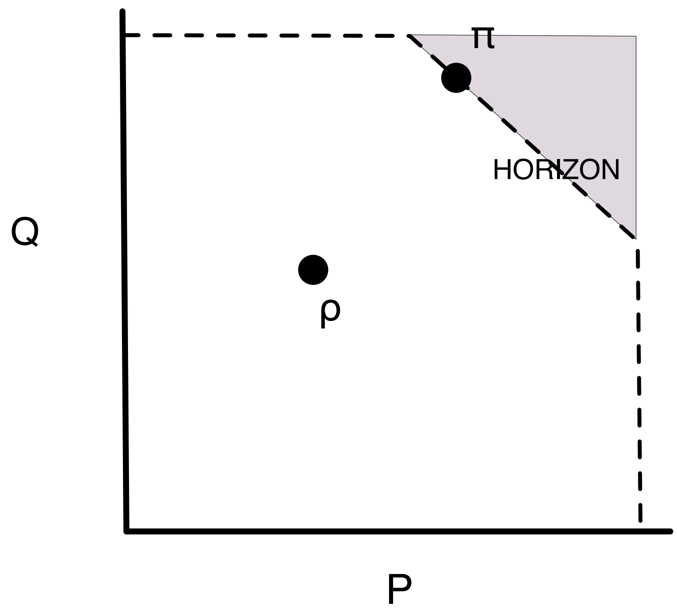


Figure 3: Illustration of a Levins tradeoff between two attributes. $\pi$ lies on the horizon of the tradeoff, illustrated by the dashed line. In this particular case, there is a region of positive value about the simultaneous in-principle maxima that the model attributes cannot reach, designated by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

Finally, two attributes exhibit a Levins tradeoff (L-tradeoff) when the magnitude of both of these attributes cannot be simultaneously maximized. If we look in $\Lambda$ 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 $\Lambda$.

**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 $\pi$ be an element of $\Lambda$. Further, let $p_{\max}$ be the maximum value for $P$ in $\Lambda$ and $q_{\max}$ be the maximum value for $Q$ in $\Lambda$. 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's philosophical work (Weisberg, 2006). Levins tradeoffs are only of significance when both model attributes in question have a maximum magnitude in $\Lambda$, otherwise a Levins tradeoff vacuously obtains.

Whenever an L-tradeoff obtains 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. Because precision and generality are often both important in scientific practice, it is likely that a theorist cannot maximize either of these attributes and will have to choose between having intermediate magnitudes for the two 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 fairly 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 justification of this entailment is more complex, so we give a formal proof below:

**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 $p_{\text{max}}$ be the maximum value for $P$ in $\Lambda$ and $q_{\text{max}}$ be the maximum value for $Q$ in $\Lambda$.

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 $\langle p_{\text{max}}, q_{\text{max}} \rangle$ is an element of $\Lambda$. Let $X$ designate this element. Therefore $(X)_1 = p_{\text{max}}$ and $(X)_2 = q_{\text{max}}$.

For a strict tradeoff to hold, there must be at least two distinct elements in $\Lambda$. Let $Y$ designate an arbitrary element of $\Lambda$ 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 $\pi$ and $Y$ for $\rho$, then $Y$ for $\pi$ and $X$ for $\rho$, 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 $(X)_2 = q_{\text{max}}$, 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 $(X)_1 = p_{\text{max}}$ it must be that $(X)_1 = (Y)_1$ □

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

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

For both to be false, it must be the case that $(X)_2 > (Y)_2$ is false, so either $(X)_2 < (Y)_2$ or $(X)_2 = (Y)_2$ must be true.

Since $(X)_2 = q_{\text{max}}$, it must be that $(X)_2 = (Y)_2$

□

$X$ and $Y$ are distinct elements of $\Lambda$, 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 attributes. 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 i-tradeoffs entail L-tradeoffs as long as a weak condition is met. If we know that an i-tradeoff holds between two attributes, and there is at least one member in $\Lambda$ where neither attribute is maximal, this entails that no member in $\Lambda$ will have both attributes at maximum. Therefore, an L-tradeoff will also 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 $\Lambda$ and $q_{\max}$be the maximum value for $Q$ in $\Lambda$. Assume that $P$ and $Q$ exhibit an i-tradeoff. Applying our definition, we know that $\forall\pi\forall\rho[((\pi)_1 < (\rho)_1 \rightarrow (\pi)_2 \geq (\rho)_2)\&((\pi)_2 < (\rho)_2 \rightarrow (\pi)_2 \geq (\rho)_2)]$.

Assume that one element in $\Lambda$ 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 $\Lambda$. 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 $[(Y)_1 < (X)_1 \rightarrow (Y)_2 \geq (X)_2]\&[(Y)_2 < (X)_2 \rightarrow (Y)_2 \geq (X)_2]$

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 $\Lambda$ 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 $\Lambda$ to certain sets of values may entail other tradeoffs, but we do not believe that any other form of attenuation is entailed by the existence of an L-tradeoff alone. However, the existence of an L-tradeoff can shape and constrain scientific practice in significant ways. In particular, the existence of an L-tradeoff means that a single or small set of models will not allow the theorist to maximize every desirable attribute. It was in this spirit that Levins argued from the existence of tradeoffs to strategies of model building, each of which sacrificed one desideratum in order to maximize others.

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.

An example of an empirically generated L-tradeoff occurs quite generally in medical diagnostic tests. Medical tests have the attributes of sensitivity and specificity. Sensitivity refers to what proportion of people with the disease are identified as such by the test. That is, a highly sensitive test will not "miss" people with the disease. Conversely, the specificity of the test reflects the proportion of people without the disease that the test correctly identifies with a negative result. A highly specific test will not mistakenly categorize disease-free people as having the illness. Note that there is nothing about the definitions of these attributes which means that they cannot be simultaneously maximal in the same test. It is entirely possible that a test might positively identify all and only those with the disease as having the disease. However, in nearly all actual cases this cannot occur. Whatever feature is used by the test in order to differentiate the groups will usually overlap at least slightly between those groups. In cases such as this, the sensitivity and specificity of the test cannot be simultaneously maximized. (See figures 4 and 5)

This is an example where two attributes cannot be simultaneously maximised due to contingent facts about the target of interest, and we believe there are many more. In this case, the tradeoff occurs because the feature used to differentiate the groups exhibits some overlap between those groups, but it is likely that L-tradeoffs occur for a number of different reasons and in different settings.

### 3.4 Summary of the tradeoffs and their interactions

From the above, we see that 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. 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 both attributes. If an i-tradeoff holds between two

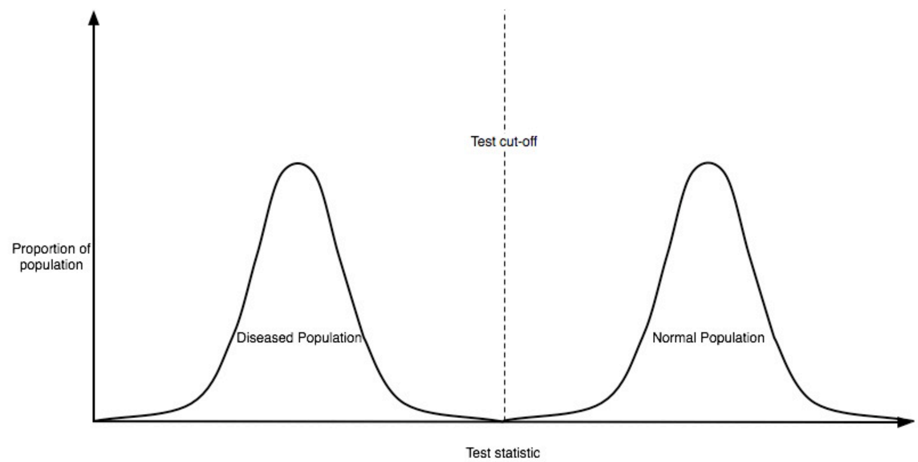


Figure 4: A situation where a diagnostic test can be both maximally specific and maximally sensitive. This occurs in any case where there is no overlap between the two populations to be distinguished relative to the test statistic.

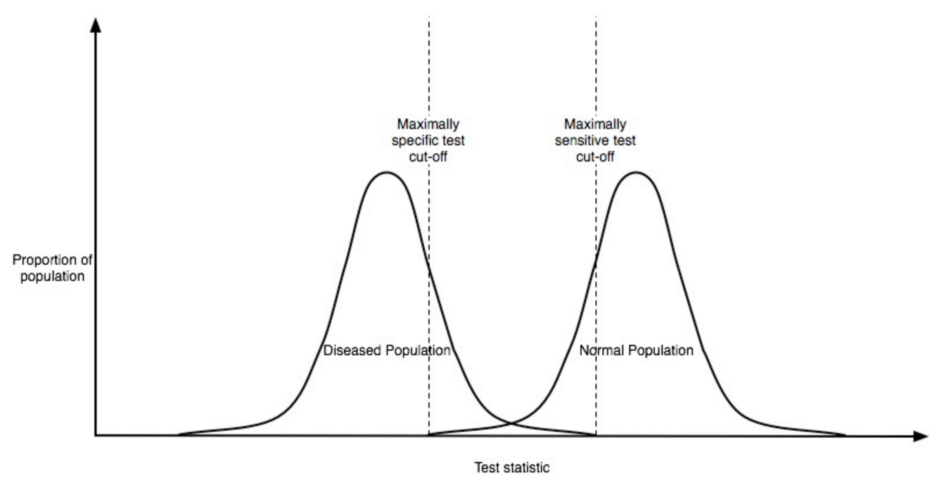


Figure 5: A situation in which a diagnostic test cannot be maximally specific and sensitive. This occurs in any case where there is an overlap between the two populations to be distinguished. When designing and calibrating the test, a decision must be made between whether a maximally sensitive *or* specific test is required, or if a test that is sub-maximal but adequate in both respects is most appropriate.

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 turn to the example of precision and generality. Philosophers of science have been especially interested in tradeoffs involving generality because of its significance in accounts of scientific method. For example, generality features prominently in many philosophical theories of scientific explanation (e.g., Strevens, 2004; Woodward, 2003), and in some accounts, generality is considered to be the core of a scientific theory's explanatory power (Kitcher, 1981, 1989; Friedman, 1974). The value of generality has also been defended by scientists, notably population biologists. These scientists point to the especially high value of general models in allowing theorists to examine disparate but similar phenomena in the same framework, exposing underlying patterns among these phenomena. (e.g., Roughgarden, 1979, 1997; May, 2001; Nowak & May, 2000)

We agree with the consensus of earlier literature and will argue that there are attenuation relationships between precision and generality. However, here we will show that some of these relationships are stronger than was previously identified, and that there are more cases of tradeoffs between precision and generality than were recognized in earlier analyses.

Before we can investigate the relationships between precision and generality, we will need to define how we will use these terms, as well as an additional feature of modeling we call *fidelity criteria*.

### 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 $\delta p$ 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 $n$-dimension 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 description. 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 of course not a simple or uncontroversial issue. This means that before we proceed, we need to to be clear regarding how we view the model-world relationship in this context.

In the semantic view of theories literature, the model-world relationship is described as one of isomorphism between the model and at least some aspects of a real-world target system. Internal critics of this picture have pointed to the stringency of this requirement and suggested other, weaker, model-theoretic relationships that may be more appropriate. (Lloyd, 1994; da Costa & French, 2003) Other critics have suggested that even these weaker correspondences do not characterize the kinds of relationships that theorists demand of their models. Some philosophers have advocated replacing a formalized model-world relationship with a more intuitive notion of similarity, either behavioral (Cartwright, 1983) or structural (Giere, 1988; Godfrey-Smith, 2006). If models are conceived of in terms of sets of trajectories in a state space, or *trajectory spaces*, then metric relationships such as the closeness of a data set to a trajectory in the trajectory space may also be a possible way to understand the model-world relationship. 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 because we want to emphasize something a bit different than most of the discussions 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 structural similarity measures, it is also important to understand the standards of fidelity applied by the modeler. We call these standards *fidelity criteria*.

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.

These differences in fidelity criteria are most apparent when we consider systems that are sensitively dependent on their initial conditions. For example, one reason we might model the traffic in Los Angeles is to anticipate and prepare for particularly bad episodes of traffic congestion. This may be achieved with a model that simply picks out patterns in past data regarding LA's traffic — the fact that congestion tends to be high between the hours of 7-10am, for example. In this case, the model will be judged according to how well it predicts future traffic jams, while whether it identifies any of the true *causes* of this congestion is of no consequence. Alternatively, we might model LA's traffic in order to help inform city planners how to minimize traffic problems when constructing a new road network. In this case, model builders will want to capture as much of the causal structure underlying LA's traffic patterns as possible. If we accept that traffic in a large city is sensitive to initial conditions, unless the initial parameter values are extremely exact, a model that includes a great deal of causal structure is unlikely to be a good predictor of future traffic patterns. However, this fact is not likely to matter to the model user in these circumstances, because they are less concerned with accurately predicting Los Angeles traffic than with representing the processes affecting it.

As we can see from these examples, there are really two kinds of fidelity criteria. One kind of fidelity concerns 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)

The fidelity criteria in use in any modeling situation will have a notable effect on generality. 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 generality we will assume that the fidelity criteria used to assess the relationship between models and target (whatever these criteria happen to be) will be held fixed.

### 4.3 Generality

Unlike precision, which is an attribute of the equations we use to describe models, generality is an attribute of the model-target relationship. Given a set of fidelity criteria, generality is a measure of how many targets the models in question apply to. There are two ways in which this concept 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 they must be analyzed separately when we consider whether they trade-off 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.[1] We call these different types of generality *a-generality* and *p-generality* respectively. P-generality is not something that only philosophers 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 are thought to help explain why there are no such

[1] We 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 to what is known to be physically impossible. Future analyses of tradeoffs might fruitfully explore more restricted modalities.

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 quite different relations to other modeling attributes, and therefore will 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. This means that we cannot always order the generality of models or sets of models according to cardinality, but will sometimes 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 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 modeling context. 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 the tradeoff between precision and generality.

### 5.1 Precision and P-generality

First we consider how an increase in precision affects p-generality. Recall that precision is a 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, in applicable cases 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 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 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 in decline. 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, the targets that models in particle physics are directed towards will often be homogeneous. In physics, it is possible to have a precise model description and still pick out a set of models that applies to most or all actual target systems. For example, a model description that targets the mass and charge of electrons can be extremely precise and still pick out a set of models that applies to all electrons. This is because physical quantities such as the mass and charge of fundamental particles are maximally 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 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 our now-familiar argument form:

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 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 then $d$.

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 we have already alluded to cases where the attenuation between them is stronger; there are general features of targets that, when present, lead us to predict a more costly interaction. In addition, changes in the scope and fidelity criteria adopted by a modeler can strengthen or weaken the tradeoff costs between precision and a-generality. Recall that fidelity criteria refer to how a modeler assesses the degree of similarity between model and target. The *scope* of a model refers to the properties of the target systems the modeler wishes to capture with that model.

When the targets of interest are heterogeneous, the attenuation between precision and generality is likely to be very costly. However, these costs can be limited if the modeler restricts her scope. For example, models with broad scope that attempt to capture the foraging behavior of all the species in a particular region of a rainforest cannot be precise without heavy costs in a-generality. But, if one is selective regarding which properties of these complex targets we wish to model, disparate targets can be made to appear more homogeneous. The scope might be restricted to the energetic aspects of the foraging, for example. Since these factors rely on biochemistry and the distribution of resources in the ecosystem, individual differences among organisms will be considerably diminished.

Another possibility is to retain the original scope, but to lower 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, often the sets of these models will become more a-general. In this way, the tradeoff between precision and model set a-generality can be made less costly.

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 Conclusions

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

As an example of these enduring constraints, we have shown that there is an s-tradeoff betwen precision and p-generality, and an i-tradeoff between precision and a-generality. Recalling the entailment relations between 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. The tradeoffs described in this paper, only a fraction of the actual tradeoffs, affect the explanatory and descriptive power of models and sets of models. Models with limited p-generality may have less explanatory power (Strevens, 2004) and models with limited a-generality can only represent a restricted set of targets. Therefore, if either explanatory power or descriptive breadth are of importance to the modeler, whenever they are faced with a heterogeneous set of targets, it is likely that the equations used to describe the models will be imprecise, or alternatively the scope or fidelity criteria employed by the modeler must be significantly limited.

More broadly, the foregoing gives us yet another 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*.

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: Univeristy of Minnesota Press.

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.

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* (Landmarks in Biology ed.). 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., & 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 approachs to explanation unified — causally. *Noûs*, *38*, 154-179.

Weisberg, M. (2004). Qualitative theory and chemical explanation. *Philosophy of Science*, *71*, 1071–1081.

Weisberg, M. (2006). Forty years of "the strategy". *Biology and Philosophy*, *21*, x–x.

Weisberg, M. (2007). Who is a modeler? *British Journal for Philosophy of Science*, *58*.

Woodward, J. (2003). *Making things happen*. Oxford: Oxford University Press.