The "Structure" of the "Strategy": Looking at the Matthewson-Weisberg Trade-off and Its Justificatory Role for the Multiple-Models Approach
Author(s): Michael Goldsby
Source: *Philosophy of Science*, Vol. 80, No. 5 (December 2013), pp. 862-873
Published by: The University of Chicago Press on behalf of the Philosophy of Science Association
Stable URL: http://www.jstor.org/stable/10.1086/673728
Accessed: 09/01/2014 10:58

The "Structure" of the "Strategy": Looking at the Matthewson-Weisberg Trade-off and Its Justificatory Role for the Multiple-Models Approach

Michael Goldsby*†

The multiple-models approach, which has its origins in Levins's work, is gaining broader acceptance among philosophers. Levins asserted that there is a trade-off between modeling desiderata, which justified the multiple-models approach through two separate justificatory paths. Some attention has been paid to the trade-off thesis, culminating in a paper by Matthewson and Weisberg. However, no attention has been paid to how the trade-off is supposed to justify the multiple-models approach. I argue that a trade-off between generality and precision cannot support one of Levins's justificatory paths, and I consider what that might mean for the multiple-models approach.

1. Introduction. One approach to modeling complex systems that seems to be gaining broader acceptance among philosophers of science is the multiple-models approach.[1] The *multiple-models approach* is the practice of using multiple (often incompatible) models to study a target system. This approach has its origin in Richard Levins's (1966) influential paper on modeling in population biology. In that paper, Levins asserts that there is a trade-off between the modeling desiderata of generality, realism, and precision such that they cannot all be maximized at the same time. He offers no argument for the the-

*To contact the author, please write to: School of Politics, Philosophy, and Public Affairs, Washington State University, PO Box 644880, Pullman, WA 99164-4880; e-mail: michael.goldsby@wsu.edu.

†I thank Elliott Sober, Malcolm Forster, Dan Hausman, Jay Odenbaugh, James Justus, John Basl, Matthew Kopec, Hayley Clatterbuck, Roberta Millstein, John Matthewson, Trevor Pearce, Roman Frigg, and Michael Weisberg for their helpful comments and advice.

1. Major proponents include Wimsatt (1980, 1987), Godfrey-Smith (2006), and Weisberg (2007a).

Philosophy of Science, 80 (December 2013) pp. 862–873. 0031-8248/2013/8005-0010$10.00

sis but instead uses it to support two further claims: (*a*) complex systems cannot be directly and adequately represented with a single model, and (*b*) no single model can meet every scientific goal that relates to a target complex system. Levins uses these latter two claims to justify the multiple-models approach.

The fact that Levins provides no argument for his trade-off thesis has attracted philosophical attention beginning with Orzack and Sober (1993), who criticize Levins for failing to do so and, most recently, in a paper by Matthewson and Weisberg (2009), wherein they demonstrate that a Levinsonian trade-off obtains between generality and precision. What has not been discussed is whether Levins's trade-off or its successor—the Matthewson-Weisberg trade-off (MWT)—provides any justification for the multiple-models approach. Although Levins suggests that there are two paths by which a trade-off might justify the multiple-models approach, in this paper, I intend to focus on the latter (*b*).

Prediction and explanation are important goals in science. Hempel (1965) famously argued for his symmetry thesis according to which every explanation is a potential prediction and every prediction can serve as an explanation. Few hold that Hempel's symmetry thesis is true; however, the mere denial of the symmetry thesis is not sufficient to justify the multiple-models approach. If, on the other hand, it could be shown that there are cases where a single model cannot meet both goals, then that might provide some justification for the multiple-models approach with respect to at least those cases. Let us call the claim that no single model can meet both our predictive and explanatory goals with respect to a particular target system the *bifurcation of scientific goals thesis* (BSG), and what we want to know is whether BSG holds for any target systems.

The main issue that I will consider in this article is whether there is any context in which BSG is true. In the next section, I will lay some groundwork and take a critical look at MWT. Then, I will consider whether the MWT provides any support for BSG, and I will argue that it does not (sec. 3); however, I will also argue that there are reasons other than MWT to think that BSG is true at least in some circumstances (sec. 4). Finally, I will conclude by briefly considering what the BSG might say about the multiple-models approach.

2. The Matthewson-Weisberg Trade-off. Matthewson and Weisberg (2009) provide an excellent analysis of the formal structure of trade-offs. In that paper they identify three types of trade-offs (from strongest to weakest): *strict trade-offs*, *increase trade-offs*, and *Levins trade-offs*. For our present purposes, the relevant trade-off is the increase trade-off. An increase trade-off occurs between two attributes, when the magnitudes of the two attributes can never be increased at the same time, but it differs from strict trade-offs insofar as it is possible for one attribute to be increased without the other decreasing. MWT is the increase trade-off that obtains between precision and generality.

Matthewson and Weisberg define *precision* as a measure of specificity of a model's parameters. In order to illustrate this concept, a couple of distinctions are necessary. Models can be instantiated or uninstantiated. An *instantiated* model is one in which the point values for all parameters are provided. An *uninstantiated* model is one that has at least one parameter that is adjustable. Consider the following models:

$$y = a + bx + cx^2, \tag{PAR}$$

$$y = a + bx, \tag{LIN}$$

$$y = 4 + (3 \pm 1)x, \tag{PSP}$$

$$y = 4 + 3x. \tag{INS}$$

In these models, y is the dependent variable and x is the independent variable, but a, b, and c are adjustable parameters, whereas 4 and 3 in INS are fixed parameters. The number (3 ± 1) in PSP is an adjustable parameter that has been partially specified, insofar as it can take any value between 2 and 4. LIN and PAR are uninstantiated models insofar as they have adjustable parameters, whereas INS is an instantiated model as all of its parameters are fixed. PSP is an uninstantiated model as one of its parameters is only partially specified. One should not confuse instantiated models with instantiations of a model. A model M_1 is an instantiation of another model M_2 iff M_1 logically entails M_2. INS is an *instantiation* of LIN, which is itself an instantiation of PAR. Finally, two models are *nested* just in case one model is an instantiation of the other.

Now we can say when one model is more precise than another if they are nested.

> (P) For any two nested models M_1 and M_2, M_1 is more *precise* than M_2 if the set of all instantiations of M_1 is a proper subset of the set of all instantiations of M_2.

Using our example models to illustrate the concept, the set of all instantiations of INS (namely, INS) is a proper subset of the set of all instantiations of PSP, and thus, INS is more precise than PSP.

Concerning *generality*, Matthewson and Weisberg define it as a measure of how many real world systems to which a model can be applied.[2] Models that apply to more real world systems will be more general than those that apply to fewer. What it means for a model to apply to a real world system is an issue that has received some philosophical attention.[3] But, Matthewson and Weisberg, following Weisberg (2007b), adopt a pluralistic view regarding the model-target relation. Specifically, they hold that it is the standards provided by the modeler, which depend upon her research interests, that determine whether a given model can be applied to its target. These standards, or *fidelity criteria*, can vary from modeler to modeler. Some might be more lax while others might be more rigorous, which, as Matthewson and Weisberg point out, may have some effect on how general a model may be (181). Being sensitive to these concerns, they maintain that when comparing the generality of two models the same set of fidelity criteria must be applied to both. Thus, we have the following metric by which we can compare the generality of two models:

> (G) A model M_1 is more *general* than a model M_2 if the instantiations of M_1 can be applied to more real world systems than the instantiations of M_2, where the same set of fidelity criteria are used to evaluate the applicability of both models.

Using G and P, an argument can now be made for MWT. Consider two nested models M_1 and M_2. It follows from P that if M_1 is more precise than M_2, then M_1 has the smaller set of instantiations. If M_1 has the smaller set of instantiations, then it is impossible for its instantiations to apply to a broader range of actual target systems than the set of instantiations of M_2. If M_1 cannot be applied to a broader range of actual target systems, then it cannot be more general. Likewise, if M_2 is the more general of the two, then it must have a broader range of instantiations than M_1 has. If it has a broader range of instantiations than M_1, then it is less precise than M_1. So at least with nested models, MWT holds.

2. Matthewson and Weisberg (and Weisberg 2006) disambiguate generality by distinguishing between possible generality (p-generality) and actual generality (a-generality). P-generality is a measure of how many possible systems to which a model can be applied, whereas the domain of a-generality is real world systems to which it can be applied. I restrict my discussion to a-generality. I do so because the MWT is about a-generality rather than p-generality.

3. Some accounts require isomorphism (van Fraassen 1980; Suppes 2002), partial isomorphism (Da Costa and French 2003), or similarity (Giere 1988) for a model to apply to the target system.

MWT fails as a general thesis, because it does not universally hold for nonnested models. For example, consider INS and PSP*:

$$y = 8 + (6 \pm 1)x. \qquad \text{(PSP*)}$$

INS is clearly more precise than PSP*, since INS is an instantiated model, and PSP* is not. Suppose that INS applies to exactly one real world system, and that PSP* applies to none. In such a case, one could increase both generality and precision by switching from PSP* to INS. This suggests that at least sometimes MWT does not hold for nonnested models. So the scope of MWT is limited to nested models.

3. Does MWT Provide Support for BSG? Matthewson and Weisberg give two reasons for thinking that their demonstration of MWT is significant. First, they claim that the demonstration answers the call made by Orzack and Sober for such.[4] Second, they argue that the trade-off has important implications for scientific methodology. Specifically they write, "there are circumstances where it would be rational for theorists to sacrifice some precision to gain generality and hence explanatory power" (189). It is this second reason that is relevant to the present discussion. Matthewson and Weisberg never explicitly claim that they think the existence of MWT supports BSG, but Weisberg (2007a) in a separate paper cites the existence of trade-offs as an important justification for the multiple-models approach. He further associates that position with Levins, which seems to suggest that since MWT is the successor to Levins's trade-off thesis, it might be that Weisberg intends for MWT to support BSG and by extension the multiple-models approach. In order for MWT to support BSG, it must be the case that generality is associated with explanatory power and precision is associated with predictive power.

3.1. Is There an Association between Generality and Explanatory Power? Matthewson and Weisberg argue that generality is important to some accounts of scientific explanation—specifically the unificationist theories of Kitcher (1989) and Friedman (1974).[5] That is because, as Kitcher (1989) puts it, "Science advances our understanding of nature by showing us how to derive different descriptions of many phenomena, using the same pat-

4. It should be noted, however, that Orzack and Sober (1993) acknowledge that a version of the Matthewson-Weisberg trade-off obtains, as they write, "A given model will be more general when it is uninstantiated than when it is instantiated" (536).

5. They also mention the hybrid causal-unificationist account of Strevens (2009), because Strevens thinks that p-generality is important for optimizing causal models. However, I set this wrinkle aside because MWT is about a-generality.

terns of derivation again and again" (432). So, if we understand a model's applying to target systems as "patterns of derivation" that provide descriptions of their target systems, then the more general model will provide more understanding than the less general model insofar as it can be applied to more target systems.

Unificationist accounts of explanation, however, are controversial, as is the claim that generality and explanation are directly associated.[6] Still, despite the worries that many may have, there is something intuitively appealing about saying that more general models have more explanatory power because they can explain more explananda. But causal theories of explanation can capture that intuition without appealing to unification. Let us suppose we have two nested causal models for the occurrence of lung cancer—S and SA.

(S) Smoking causes lung cancer.

(SA) Smoking and exposure to asbestos each cause lung cancer.

Further suppose that we have three targets (all of whom have lung cancer): Alex who smokes but is never exposed to asbestos; Bobbi who smokes and was exposed to asbestos; and Charlie who was exposed just to asbestos. SA is more general than S in this example because it applies to all three targets, whereas S applies to only Alex and Bobbi.

There are two questions that we need to separate. (1) Which of these models is more explanatory in the sense of being able to explain more explananda? And (2) which of these models offers the best explanation for a given explanandum? Looking at how the unification and causal theories of explanation answer these questions will help clear up the relation between generality and explanation. With regard to 1, both theories would agree that SA is more explanatory, but for different reasons. According to the unification theory, SA would be the better (or only) explanatory model in virtue of its ability to unify the explananda—that is, its generality. A causal theorist would also judge SA as the better explanatory model, but in virtue of its causal inclusiveness rather than its generality. For the causal theorist, generality is a by-product of the model's causal inclusiveness. Turning to question 2, both theories would judge that SA is the better explanation for Bobbi's and Charlie's lung cancer, which seems right. However, when we take Alex's lung cancer as the explanandum, the causal theorist could say that both S and SA are adequate; whereas the unificationist is committed to say-

6. See Woodward (2003) for his objections to unificationist accounts of explanation and Hitchcock and Woodward (2003) for their critique of generality's explanatory power. See also Sober (1999).

ing that SA is the better (or only) explanation. To say that SA is a better explanation for Alex's lung cancer seems bizarre. It is hard to see why citing asbestos exposure gives us any more understanding as to why Alex has lung cancer, given that she has never been exposed to asbestos. Returning to the case of Bobbi, it seems equally counterintuitive to say that the reason that SA is a better explanation for her having lung cancer is that SA is more general. It is better to say that the reason SA is a better explanation is because it cites both of the causes that actually lead to her unfortunate situation.

So while I agree that there is a sense in which those models that are able to explain more explananda are more explanatory, I wish to deny that it is in virtue of the model's generality. An interesting consequence of Matthewson and Weisberg's definition of generality (G) is that with causal models, the more causes that the model cites make it more general. That is to say that it is the causal inclusiveness that makes the model both more general and explanatory. General models then can be more explanatory, but the reason for their increased explanatory power is due to their causal inclusiveness.

3.2. Is There an Association between Precision and Predictive Power? That there is some connection between generality and explanation seems plausible—whether it is due to greater unification or citing more causes that affect the target system. However, it is unclear as to what exactly is being sacrificed when we are slighting precision to increase generality apart from precision itself. Let us grant that increasing generality is associated with an increase in the model's explanatory power. Then, as a result of MWT such an increase in explanatory power comes at the expense of precision. That means that in order for MWT to support BSG, as Levins intended his trade-off to do, increases in precision must be associated with increases of predictive power.

One possible approach to showing that precision and predictive power are associated would be to deny that sufficiently imprecise models are capable of making predictions. The argument would run something like this. A maximally precise model—meaning one that has all of its parameters fixed—can make a point prediction about the value of the model's dependent variable once the values of all the independent variables are set. However, more imprecise models—those that contain one or more adjustable parameters—cannot make the same sort of predictions, even when the values of independent variables have been specified. That is because if any of the parameters are unspecified, then those parameters can take an infinite range of values. Since imprecise models cannot make point predictions, they cannot make testable predictions, and therefore what is sacrificed when one sacrifices precision is the model's very ability to make testable predictions.

The first problem with this line of argument is that it is simply false that in order for a prediction to be testable it must be a point prediction. Consider LIN. Even when the parameters are left unspecified, there are plenty of data sets that are inconsistent with LIN. The second problem with this line of argument is that it seems to ignore the impact of the Duhem-Quine thesis. According to that thesis, we should expect that without auxiliary assumptions an imprecise model cannot make point predictions. However, with the addition of some assumptions concerning the likely value of those parameters, the model would be capable of making predictions that are every bit as testable as those made by the maximally precise model. In fact, this seems to be exactly how scientists treat general yet imprecise models. They take the model that is supposed to apply to a wide range of target systems and estimate what the parameters would be when applied to the system of interest. Then they use the model with the assumptions about what values they think the parameters will take to make a prediction about that system. With some assumptions about the values of the parameters, the more general model can make any prediction that its more precise cousin can make, provided that the two models are nested.

Another way in which predictive power may be sacrificed is through a loss of predictive accuracy. Predictive accuracy is a measure of how well a model predicts new data when fitted to old (Forster and Sober 1994). Is it reasonable to expect that with two nested models one of which is maximally precise and the other uninstantiated that the maximally precise model will be more predictively accurate than the uninstantiated one on average? The answer is no provided that the two models have an equal number of nonzero parameters. It might be the case that an instantiated and thus maximally precise model will fare much worse than the more imprecise models if the specified parameters of the instantiated model are wildly unrealistic. My point here is that mere precision has little bearing on predictive accuracy.

4. Is There Any Support for BSG? I think that there is some reason to think that in a great number of cases—perhaps even the majority of cases involving complex systems—our best predictive model cannot be the same as our best explanatory model. My task, in what follows, is to show that it is possible for the goals of prediction and explanation to bifurcate such that if that sort of situation were to obtain, then no single model could be our best explanatory and our best predictive model. In order to do so, I will use the Akaike framework from model selection theory.

Akaike's theorem provides a framework from which we can compare the estimated predictive accuracy of two or more models. This framework uses the Akaike information criterion, also known as AIC:

The *AIC score of a model M*, $\mathrm{AIC}(M) =_{\mathrm{def}} \log\{\Pr[\mathrm{data}|\mathrm{L}(M)]\} - k$.

The term $\mathrm{L}(M)$ refers to the instantiation of M that has the highest likelihood of all the instantiations of M when compared to the data, and $\Pr[\mathrm{data}|\mathrm{L}(M)]$ is that likelihood. To determine the $\mathrm{AIC}(M)$ one takes the natural logarithm (base e) of the likelihood of $\mathrm{L}(M)$ and subtracts the number of adjustable parameters, represented by k. The log likelihood of $\mathrm{L}(M)$ is a measure of how well the model fits the data, and k represents a correction for complexity. What is important about $\mathrm{AIC}(M)$ is that it gives us an unbiased estimate of M's predictive accuracy (Forster and Sober 1994; Forster 2000, 2001; Sober 2008). The absolute value of the score is not important; what is important is that we can use AIC to compare the estimated predictive accuracy of uninstantiated models, such that we would expect that the model with the higher AIC score would be better at predicting new data.

To show how BSG might be true in some situations, I am going to tell a little fable. Suppose, for simplicity's sake, that we have recently discovered a causally isolated system, the behavior of which is causally influenced by exactly two causes. It would not be long before we would begin to study this new and exciting system. Of course, there would be some speculation as to the best way to model the system, but just as a rough set of hypotheses, we settle on three models:

(1-cause) $y = ax_1 + e$,

(2-cause) $y = ax_1 + bx_2 + e$,

(3-cause) $y = ax_1 + bx_2 + cx_3 + e$.

In these models, y represents the measurable quantity that we wish to predict, and let us stipulate that x_1 and x_2 represent the two causes that actually affect the behavior of the system. Although x_3 is some measurable quantity that is hypothesized by the 3-cause model to be causally related to the system, unbeknownst to us it is not. Of course, a, b, c, and e are adjustable parameters in the model where e is an error term that is supposed to correct for measurement errors.

As a brief aside, let us say that the explanandum in our fable is the behavior of this particular system. According to the stipulations of this example, if we take a causal view (Woodward 2003; Strevens 2009) of explanation, then 2-cause is the best explanatory model, since it cites all and only those causes that affect the behavior of the system. Of course, this last bit is not something that is epistemically available to us.

Returning to our fable, which model would we expect to be the best predictor? The answer to that question depends upon how large the data set is from which we estimate the parameters. Imagine that when we first encounter the system, we start making some observations—measuring this and that—which we use as our initial data set. Since this is just the beginning of our investigation, our data set will be quite small. From this data set we estimate the parameters of our three models. It is reasonable to expect that more complex models will fit the data at least a little bit better than simpler models, such that

$$\Pr[\text{data}|L(\text{3-cause})] > \Pr[\text{data}|L(\text{2-cause})] > \Pr[\text{data}|L(\text{1-cause})].$$

However, since the data set is small, the difference in likelihoods will also be small between the three models. As a result, the following inequality (1) would hold because the penalty for complexity would overwhelm the difference in fit:

$$\text{AIC(1-cause)} > \text{AIC(2-cause)} > \text{AIC(3-cause)}. \tag{1}$$

If (1) were true, then we would expect 1-cause to be the better predictor. At this stage in our investigation, BSG would be true, because 1-cause would be the best predictive model, while 2-cause would be our best explanatory model.

Now imagine that we go out and test our models, and in the process of doing so we collect more data. It might even be the case that we would split into two or three camps, each in favor of one of the models. The 1-causers would cite their better predictive accuracy as evidence that their model is the best, saying that only prediction matters. The 2-causers, on the other hand, might still insist that while the 1-cause model is the better predictor, the 2-cause model is still better, as it better explains the behavior of the system. However, one 2-causer proposes an interesting gambit. She suggests that the parameters should be reestimated using the now larger data set. She also proposes that the matter be settled by the most ambitious testing ever conducted on the system, such that the models will be compared to see how well they predict a new data set that is equal in size to our current larger data set.

Since the data set is now larger and the testing more ambitious, it is entirely possible that the new AIC scores would be such that AIC(2-cause) > AIC(1-cause) > AIC(3-cause). This is because the difference in fit to data between the 2-cause and 1-cause models would be enough to offset the penalty for complexity. In such a case, we would expect the 2-cause model to be the better predictor. At this point in our fable, BSG does not hold because 2-cause is both our best predictive model and our best explanatory model. Of course, a 3-causer might take her cue from the 2-causer's gambit, convincing us to go out and collect even more data and to conduct even more ambitious testing, such that our old data set grows to immense proportions, as does the testing. At some point it may even be the case that AIC(3-cause) > AIC(2-cause). This might happen because the errors in our measurement begin to look more like another cause, as our data set grows beyond proportion. It may even be the case that 3-cause at this point would be our best predictor and BSG would hold again.

5. Conclusion. What's the moral of our fable? The first is that is it possible for our explanatory and predictive goals to bifurcate. The second moral is that BSG holds only in certain situations. That is to say that there is a certain sweet spot where BSG does not hold. BSG is likely to hold in situations where our data set is meager and the causal interactions of the system of interest are many. In our real and messier world, this might more often be the case than not. I think in most cases we are on the low side of the sweet spot. It is in principle possible for us to be on the high side of the sweet spot, but that seems a bit more unlikely.

Unfortunately, the third moral is that we may not know when BSG holds or when it does not. It is unclear to me just what sort of evidence would be required to show that we are in (or out) of the sweet spot. This has some implications for the multiple-models approach. When BSG holds it does justify the multiple-models approach, but we cannot know with certainty when it does. So, rather than justifying a multiple-models approach, it might actually support a move toward instrumentalism. I think that in many cases, given the complex systems of the complex world we live in, a plausible argument can be made that we are not in the sweet spot and that BSG holds. But whether that is sufficient to provide the justification for the multiple-models approach over instrumentalism is another matter—and it is one that is worthy of future research.

REFERENCES

Da Costa, Newton, and Steven French. 2003. *Science and Partial Truth*. Oxford: Oxford University Press.

Forster, Malcolm. 2000. "Key Concepts in Models Selection: Performance and Generalizability." *Journal of Mathematical Psychology* 44:205–31.

———. 2001. "The New Science of Simplicity." In *Simplicity, Inference, and Modelling*, ed. Arnold Zellner, H. A. Keuzenkamp, and Michael McAleer, 83–119. Cambridge: Cambridge University Press.

Forster, Malcolm, and Elliott Sober. 1994. "How to Tell When Simpler, More Unified, or Less *Ad Hoc* Theories Will Provide More Accurate Predictions." *British Journal for the Philosophy of Science* 45:1–35.

Freidman, Michael. 1974. "Explanation and Scientific Understanding." *Journal of Philosophy* 71 (1): 5–19.

Giere, Ronald. 1988. *Explaining Science: A Cognitive Approach*. Chicago: University of Chicago Press.

Godfrey-Smith, Peter. 2006. "The Strategy of Model Based Science." *Biology and Philosophy* 21:625–40.

Hempel, Carl. 1965. *Aspects of Scientific Explanation*. New York: Free Press.

Hitchcock, Christopher, and James Woodward. 2003. "Explanatory Generalizations." Pt. 2, "Plumbing Explanatory Depth." *Noûs* 37 (2): 181–99.

Kitcher, Philip. 1989. "Explanatory Unification and the Causal Structure of the World." In *Scientific Explanation*, ed. Philip Kitcher and Wesley Salmon, 410–505. Minneapolis: University of Minnesota Press.

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

Matthewson, John, and Michael Weisberg. 2009. "The Structure of Tradeoffs in Model Building." *Synthese* 170 (1): 169–90.

Orzack, Steven H., and Elliott Sober. 1993. "A Critical Assessment of Levins's 'The Strategy of Model Building in Population Biology (1966).'" *Quarterly Review of Biology* 68:533–46.

Sober, Elliott. 1999. "The Multiple Realizability Argument against Reductionism." *Philosophy of Science* 66:542–64.

———. 2008. *Evidence and Evolution*. Cambridge: Cambridge University Press.

Strevens, Michael. 2009. *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press.

Suppes, Patrick. 2002. *Representation and Invariance of Scientific Structures*. Stanford, CA: CSLI Publications.

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

Weisberg, Michael. 2006. "Forty Years of 'The Strategy': Levins on Model Building and Idealization." *Biology and Philosophy* 21 (5): 623–45.

———. 2007a. "Three Kinds of Idealization." *Journal of Philosophy* 104 (12): 639–59.

———. 2007b. "Who Is a Modeler?" *British Journal for Philosophy of Science* 58:207–33.

Wimsatt, William. 1980. "Randomness and Perceived Randomness in Evolutionary Biology." *Synthese* 43:287–329.

———. 1987. "False Models as Means to Truer Theories." In *Neutral Models in Biology*, ed. Matthew Nitecki and Antoni Hoffman, 23–55. London: Oxford University Press.

Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.