Model-Building in Philosophy

Timothy Williamson

[to appear in Russell Blackford and Damien Broderick (eds.), *Philosophy's Future: The Problem of Philosophical Progress*. Oxford: Wiley]

One notable form of progress in the natural and social sciences over the past century has been the development of better and better *models* of the phenomena they study. The models are typically presented in mathematical terms: for instance, by differential equations for the rise and fall in population of a predator species and a prey species, interacting only with each other, or by a set of ordered pairs for the networking relations in a society.

When a system resists direct study, because it is so complex or hard to observe, model-building constitutes a key fall-back strategy. Studying a model often yields insight into the phenomena it models. When one model is replaced by another that captures more about how the phenomena work, science progresses.

Sometimes such progress is a step towards discovering universal laws of nature, non-accidentally exceptionless generalizations. However, macroscopic phenomena are typically too complex and messy to obey many informative exceptionless generalizations framed in macroscopic terms. (That goes for some microscopic phenomena too.) In such cases, the discovery of universal laws may not be a reasonable aim for those branches of science, even if there are still useful rules of thumb. Then it may be more realistic and more fruitful to aim at building increasingly good models instead. Special sciences such as economics and psychology are salient examples. Even in evolutionary biology, progress may consist more in the development of better models than in the discovery of universal laws.

This chapter argues that in philosophy, too, one form of progress is the development of better and better models—especially, but not exclusively, in those branches of philosophy, such as ethics, epistemology, and philosophy of language, which deal primarily with the human world in all its complexity and mess. Not only *can* philosophy make progress through model-building, it has been doing so for quite some time. Philosophers tend to feel embarrassed by the question "So what has philosophy discovered recently?" When we try to think of an informative generalization whose universal truth has recently come to be known through the efforts of philosophers, we may well not come up with much. We tend to assume that most of the natural and social sciences are doing far better. Certainly they are indeed making progress, but this may consist much less than we suppose in the discovery of universal generalizations and much more in the development of better models. Once we look for progress of that kind in philosophy, it is not hard to find. It is there right under our noses.

What are models?

Philosophers of science use the word "model" in a confusing variety of ways, as do scientists themselves. Clarity has not been served by a universalizing tendency in the philosophy of science to define the word in a way meant to apply to *all* scientific theories or *all* uses of the word "model" in science. For present purposes, a more helpful recent trend in the philosophy of science has been to use the term "model-building" to identify a specific recognizable type of theoretical activity that some but not all scientists engage in, some but not all of the time (Godfrey-Smith 2006a; Weisberg 2007). A scientific research group may advertise a position as a "modeler"; some but not all members of the group will be modelers. Similarly, I do *not* suggest that all philosophizing is model-building. Rather, some but not all philosophers build models, some but not all of the time.

Even in the restricted use of "model," there are different accounts of what models are. For the sake of simplicity and clarity, I will give my own account, but what I say could be adapted to other accounts: it is easier to agree on *whether* a scientist is presenting a model than on what sort of thing that model is.

Here, a model of something is a *hypothetical example* of it. Thus a model of predator-prey interaction is a hypothetical example of predator-prey interaction. The point of the qualification "hypothetical" is that the example is presented by an explicit description in general terms, rather than by pointing to an actual case. For instance, one writes down differential equations for the changing population sizes of the two species, rather than saying "the changing numbers of foxes and rabbits in Victorian Sussex." The description picks out a *type* of case, rather than one particular case: for instance, the type of any predator-prey interaction that obeys the given differential equations.

For the model-building methodology to work well, the description of the hypothetical example must be precise and specific enough to be formally tractable. That is, it should enable us to derive answers to many relevant questions about the example. When we explore the model, we do so on the basis of what follows from the description itself, which is designed to facilitate that process. We do not assume that the model fits the knowledge we already have of the phenomenon under study, since that is one of the main questions at issue. But if the fit turns out to be reasonably good, exploring the model becomes a way of indirectly exploring the original phenomenon. The mathematical clarity of the description helps make direct study of the model easier than direct study of the phenomenon itself.

The hypothetical example, the type picked out by the description, may or may not have actual instances. Indeed, it may or may not have *possible* instances. For example, evolutionary biology typically uses differential equations for population change, even though they treat the change in the number of group members as continuous whereas really it must be discrete; answers to "How many?" questions do not form a continuum. Strictly speaking, such a model is *impossible*; it is a type metaphysically incapable of having instances. But that does not mean that the model collapses. The differential equations are mathematically consistent; we can still make a stable tripartite distinction between what follows from them, what is inconsistent with them, and what is neither. Moreover, the mathematical consequences of the description may turn out to be similar enough to descriptions in similar terms of the observed behaviour of the target real-life phenomenon for the model to provide considerable theoretical insight into the target. In advance, we might not have expected impossible models to have such cognitive value, but it has become clear that they can.

The role of formal consistency in a model-building methodology provides a link between this meaning of "model" and its meaning in mathematical logic. In the logical sense, a model of a theory (call it a "logic-model") is an interpretation of the theory on which it comes out true. The interpretation must give the purely logical expressions (such as "if" and "not") their intended interpretations but may radically reinterpret non-logical expressions (for instance, by treating the word "fox" as applying to numbers). A theory is *logically consistent* if and only if it is true on at least one such interpretation, in other words, it has a logic-model. A sentence *logically follows* from a theory if and only if it is true on every interpretation on which the theory is true, in other words, every logic-model of the theory is a logic-model of the sentence.

We can apply those logical distinctions to model-building in science by treating the description of the model as a mini-theory, and the purely mathematical expressions in the description as logical, so that their interpretation is held fixed. On its intended interpretation, the description of the model may pick out an impossible type (for instance, because it describes population growth as continuous). Nevertheless, the description is logically consistent, because it has a logic-model: it is true on an unintended interpretation.

The mathematical clarity of the description typically makes such a logic-model easy to construct, for instance by reinterpreting its non-logical terms (such as "predator" and "prey") as applying to purely mathematical entities with the right formal structure. The purely logical consequences of the description do not depend on the intended interpretations of its non-logical terms; they are determined by the whole class of intended and unintended interpretations alike. Nevertheless, the non-logical terms are not idle, for they are needed to co-ordinate comparisons with the real-life phenomenon. If we interchange the words "predator" and "prey" in the description, the comparisons go differently.

The simplified and sometimes idealized nature of models is no surprise on this account. They are typically intended to be easier to explore than the real thing; simplicity and idealization contribute to that.

A warning is in order. The talk of *building* models might suggest a constructivist philosophy of science, on which model-building is a matter of invention rather than discovery, and is not in the business of uncovering truths independent of the inquiry itself. But that would be a very naïve conclusion to draw. Rates of population change in predators and prey are not figments of the scientific imagination. If we are investigating a complex reality out there, it is not at all surprising that it is sometimes best to use a sophisticated, indirect strategy, to ask questions quite subtly related to the overall aims of the inquiry. To build a model is just to identify by description a hypothetical example which we intend to learn about in hope of thereby learning about the more general subject matter it exemplifies. Nothing in that strategy is incompatible with a full-bloodedly realist nature for the scientific inquiry. The same goes for model-building in philosophy.

On a full-bloodedly realist conception of model-building, we should expect it under favourable conditions to provide *knowledge*. But, since only what is true is known, and virtually no model description is strictly true of its real-life target, what knowledge can model-building provide? What could its content be?

When we explore a model by valid deductive reasoning from the model description, we learn necessary truths of the general conditional form "If a given case satisfies the model description, then it satisfies this other description too." That broadly logico-mathematical knowledge has the virtue of precision, but by itself is less than we want, since it says nothing unconditional about how close the original phenomenon (such as predator-prey interaction) comes to satisfying the model description. Fortunately, we can also learn unconditional though vaguer truths of the general form "This model description fits the phenomenon better than that one does in the following ways," where the fit is usually approximate. Although much more needs to be said about what such approximation consists in, for present purposes the general picture will do. Such a combination of precise conditional knowledge and vague unconditional knowledge of the target is ample reward for the work of model-building. (Weisberg 2013 gives a far more detailed account of model-building in science.)

Models in philosophy

The need for model-building is hardest to avoid where the complex, messy nature of the subject matter tends to preclude informative exceptionless universal generalizations. The paradigm of such complexity and mess is the human world. Hence the obvious places to look for model-building in philosophy are those branches most distinctively concerned with human phenomena, such as ethics, epistemology, and philosophy of language. Of course, categories like goodness and duty, knowledge and justification, meaning and communication are not restricted to humans. Even those that do not apply to non-human animals on earth can in principle apply to actual or possible non-human agents, perhaps vastly more sophisticated intellectually than we will ever be. Philosophers typically want their theories to apply to such non-human agents too. But that only makes exceptionless universal generalizations still harder to find. By contrast, pure logic supplies fertile ground for powerful exceptionless universal generalizations. One might expect the same of fundamental metaphysics too. Although the metaphysical question of personal identity looks more complex and messy, it also looks less fundamental.

As it happens, the few extant discussions of model-building in philosophy have tended to concentrate on model-building in metaphysics (Godfrey-Smith 2006b, 2012; Paul 2012). One reason is perhaps that metaphysics has the worst press of any branch of philosophy, so the need for a new methodological defense may be felt most strongly there. Model-building is indeed sometimes used even in fundamental metaphysics. An example is the idea of *gunk*, stuff (or space itself) of which every part has a lesser part, so it has no perfectly atomic parts. Gunk may not be actual, but is it metaphysically possible? It is very tricky to work out which natural assumptions about the part-whole relation are logically consistent with gunk. Constructing mathematical models of gunk provides a good way of answering such questions (see Arntzenius 2008, Russell 2008, and Wilson 2008 for a debate).

If we turn to more obviously likely branches of philosophy, such as epistemology and philosophy of language, examples of model-building are easy to find.

In epistemology, a standard model of epistemic uncertainty is a lottery. Here is a typical description:

There are exactly 1000 tickets in the lottery, numbered from 1 to 1000. Exactly one will win. The lottery is fair. That is all you know about it. Thus, on your evidence, each ticket has probability 1/1000 of winning.

That description involves various assumptions typically false of lotteries in real life. For instance, it assumes that it is certain on your evidence exactly how many tickets will be in the draw. Nevertheless, a good test of epistemological theories is to work out what they say about this simple case. For instance, consider the proposal that you should accept a proposition if and only if it is at least 90% probable on your evidence. If so, you should accept that the winning number will be greater than 100, and you should accept that it will be at most 900, but it is not the case that you should accept that it will be greater than 100 and at most 900. You are obliged to accept one conjunct and you are obliged to accept the other, but you are not obliged to accept their conjunction. That is at best an uneasy combination. One can show that a similar problem arises for any probabilistic threshold for acceptance more than 0% and less than 100% (varying the number of tickets when necessary). Although lottery models are elementary, they already have enough structure to make trouble for many superficially attractive ways of thinking about uncertainty. Moreover, their simple mathematical structure makes it trivial to define mathematical logic-models with that structure, so their consistency is not in doubt.

The branch of epistemology known as *formal epistemology* is much concerned with model-building. The models come from two main sources. Some, like that above, are probabilistic, often in the Bayesian tradition of thinking about probability, which has been hugely influential in the natural and social sciences (for example, Howson and Urbach 1993). Others are models associated with epistemic logic in a rich tradition originating with Jaakko Hintikka (Hintikka 1962; Ditmarsch, Halpern, Hoek, and Kooi 2015): although not all standard logic-models of epistemic logic are models in the present sense of epistemic situations, they can all be reinterpreted in a natural way as such models. One can also add probabilities to models of epistemic logic in a natural way

(Williamson 2000). Model-building in epistemic logic has found numerous applications in computer science and theoretical economics, for instance in understanding the relations between public and private knowledge. When one looks back on the vast body of results produced by model-building in formal epistemology over the past half-century, it seems idle to deny that considerable progress has been made in understanding the epistemic subtleties of many kinds of situation. Nor should one imagine that the progress is primarily mathematical. Although mathematics is usually involved, as in model-building throughout the natural and social sciences, the main interest of the models is not in their abstract mathematical structure but in their epistemic interpretation.

In the natural and social sciences, models are often tested by their predictions of measurable quantities. Models of epistemic logic typically make no such predictions, so how are they to be tested? But even in the natural and social sciences, models are often tested by their qualitative predictions (Weisberg 2013: 136). Models of epistemic logic can be tested in that way too. For instance, in a common type of model for epistemic logic, whenever one fails to know something, one knows that one fails to know it. There are many counterexamples to that principle. For instance, a flat-earther fails to know that the earth is flat, because it isn't flat, but he also fails to know that he fails to know that it's flat, because he thinks he *does* know that it's flat. For some purposes we can legitimately abstract away from such cases. But once we become interested in the limitations of self-knowledge, such cases matter, and our models must permit them. Of course, such qualitative testing presumes that we have some model-independent knowledge of the target phenomenon, but that is equally true of quantitative testing. If we started in total ignorance of the target, we could hardly expect to learn much about it by modelling alone.

Many developments in philosophy of language can also be understood in model-building terms.

Originally, Frege and Russell introduced formal languages into philosophy as languages in which to carry out proofs more rigorously than was possible in natural languages, because the formal languages were more precise and perspicuous. That was not model-building. Later, Russell and the younger Wittgenstein argued that such formal languages articulated the covert underlying structure of ordinary thought and language. That was still not model-building.

Carnap did something different. He defined the syntax and semantics of simple, artificial examples of languages in meticulously explicit detail (Carnap 1947). He did not intend to work *in* these languages, nor did he intend them to have the expressive power of natural languages. Rather, he intended them as *models of language*, to show exactly how his intensional semantics could in principle assign meanings to all the expressions of a language. It did so *compositionally*, determining the meaning of a complex expression as a function of the meanings of its constituents, in a way that explains how we can understand new sentences we have never previously encountered by understanding the familiar old words of which they are composed and the ways in which they are put together.

The key challenge was to explain how modal operators like "possibly" and "necessarily" work. They did not fit the available model for sentence operators, truth-functionality. Operators like "and," "or," and "not" are truth-functional in the sense that they are used to form complex sentences out of simpler ones, where the truth-value of the former is determined by the truth-values of the latter. For instance, the conjunction "A and B" is formed from the simpler sentences A and B; it is true if they are both true, false if one of them is false. But modal operators are not truth-functional. That the sentence A is false does not determine the truth-value of "Possibly A," which depends on whether A is contingently false or necessarily false.

Carnap solved the problem by taking as the crucial semantic property of a sentence not its *extension*, its actual truth-value, but its *intension*, its spectrum of truth-values across all possible worlds (in his terminology, "state-descriptions"). Although the extension of A does not determine the extension of "Possibly A," the intension of A *does* determine the intension of "Possibly A." For if the intension of A has truth at some possible world, then the intension of "Possibly A" has truth at *every* possible world, while if the intension of A has truth at no possible world, then the intension of "Possibly A" also has truth at no possible world.

Carnap's insight is the root of the immensely fruitful tradition of possible worlds semantics, which has been central to later developments in both philosophy of language and formal semantics as a branch of linguistics. Although various aspects of his account are no longer widely accepted, it still constitutes major progress. He provided a simple working model of the semantics of a language with modal operators. Much subsequent work in formal semantics has in effect provided increasingly sophisticated model languages whose expressive power comes increasingly close to that of natural languages. Even if one thinks that formal models can never capture all the untidy complexity of natural languages, it is obscurantist to conclude that they provide no insight into the workings of natural languages, just as it would be obscurantist to claim that formal models in natural science provide no insight into the untidy complexity of the natural world. (One might even treat the later Wittgenstein's carefully described language games as partial models of language, emphasizing links to action and imperative rather than indicative utterances, intended as a corrective to over-emphasis on language's descriptive function. Presumably, he would have hated their assimilation to a scientific method.) The philosophical significance of those semantic insights extends beyond philosophy of language. For instance, philosophers in virtually all branches of the subject ask what is possible or necessary. If they use such modal terms in their reasoning with no reflective understanding of how their meanings work, they are liable to commit logical blunders.

The future may well see radical changes in the overall theoretical frameworks within which epistemic, semantic, and other models are built. Nevertheless, it is reasonable to expect that insights embodied in current models will be preserved, refined, and deepened in models constructed within those future frameworks, just as happens in the natural and social sciences.

Perhaps, in the future, research groups in philosophy will advertise positions for modelers.

Methodological reflections

Not all the advantages of formal methods in philosophy depend on model-building. Sometimes one formalizes the premises and conclusion of a tricky philosophical argument in order to show that the latter follows from the former in a recognized proof system for the formal language. That is progress, but it is not model-building in any distinctive sense.

Model-building is more relevant to showing that a conclusion does *not* follow from some premises. As section 1 noted, model descriptions facilitate the construction of uncontentious logic-models with the appropriate mathematical structure. When a model description seems informally consistent with the premises but not with the conclusion of a philosophical argument, one can often construct a corresponding logic-model on which the premises are true but the conclusion false, and thereby demonstrate that the conclusion does not logically follow from the premises. As a special case, when a model description seems informally consistent with a philosophical theory, one can often construct a corresponding logic-model on which the theory is true, and thereby demonstrate that it is logically consistent: it does not logically entail a contradiction.

Of course, those logical relations are not all that matters; a logically consistent theory may still be obviously false, and a conclusion that does not follow logically from some premises alone may follow from their combination with some obvious truths as auxiliary premises. But the same model-building methodology helps us track those further logical relations too. Thus one advantage of model-building—not the only one—is to make us more efficient and accurate at mapping the logical space in which we are theorizing. Without such a map, we blunder about in a fog, bumping into unexpected obstacles, falling over cliffs. It is not at all uncommon for elaborate philosophical theories to suffer some form of logical collapse: if not inconsistency, the erasing of vital distinctions. Many such disasters could have been avoided if the theory's proponents had thought to subject it to preliminary testing by model-building, for instance by trying to build a model yielding a non-trivial logic-model on which the theory came out true.

For the efficient mapping of logical relations, the advantages of simple models are obvious. Simplicity conduces to computational feasibility, so that we can in practice derive the model's mathematical properties by deductive reasoning from its description. This is particularly important for the strategy of learning about the target phenomenon by manipulating the model, adjusting it (by varying the values of parameters or in other ways) to see what difference it makes—for instance, whether a promising feature of the model is robust under such perturbations. One can gain large cognitive rewards, as well as pleasure, from playing even with a toy model, because such variations are so easy to track.

Simple models have other, less obvious advantages. One is the avoidance of arbitrary features. The more adjustable parts a model has, the more opportunities it offers the model-builder to rig the results, to gerrymander the model by setting parameters and arranging structure in *ad hoc* ways to fit preconceived prejudices. Simplicity, elegance, symmetry, naturalness, and similar virtues are indications that the results have not been so rigged. Such virtues may thus ease us into making unexpected discoveries and alert us to our errors.

Simplicity is often connected with idealization. An idealized surface is frictionless; an idealized planet is a mass at a point. Those idealizations simplify the mathematics. But idealization is also a means of abstracting from "noise," complicating factors that interfere with, and obscure, the phenomenon we are trying to understand.

Here is an instance from formal epistemology. Standard epistemic logic treats agents as *logically omniscient*: the structure of its models presupposes that if one knows some propositions, one also knows any other proposition they entail. Standard probability theory makes a similar though slightly weaker assumption: if one proposition entails another, the latter is at least as probable as the former. Such models ignore the computational limits of actual agents. Even if two mathematical formulas are logically equivalent, we may accept one but not the other because we are unaware of their equivalence; mathematics is difficult. However, idealizing away such computational limits is not just a convenient over-simplification. One may be interested in the epistemological effects of our *perceptual* limits: our eyesight is imperfect; our powers of discrimination by sight are limited. Since ignorance may result from either perceptual or computational limits, we must separate the two effects. A good way to do that is by studying models where the agent resembles a short-sighted ideal logician, with perceptual limits but no computational limits, whose ignorance therefore derives only from the former. For that purpose, the structure of standard models of epistemic logic is just right (Williamson 2014). More generally, model-building allows us to isolate one factor from others that in practice always accompany it.

Although model-building already plays a significant role in philosophy, philosophers have not fully adjusted to its methodological implications. For instance, *counterexamples* play a much smaller role in a model-building enterprise than they do in traditional philosophy. The traditional philosopher's instinct is to provide counterexamples to refute the simplifications and idealizations built into a model, which rather misses the point of the exercise. A theoretical economist once remarked to me that a paper like Gettier's classic refutation of the analysis of knowledge as justified true belief by means of a couple of counterexamples (1963) would be considered unpublishable in economics. For economics is primarily a model-building discipline: since no model is expected to fit the actual phenomena perfectly, pointing out that one fails to do so is not considered newsworthy. What defeats a model is not a counterexample but a *better model*, one that retains its predecessor's successes while adding some more of its own. For reasons explained at the end of section 1, that does not mean that model-building disciplines are unconcerned with truth. They too pursue truth, but by more indirect strategies. Of course, it is unfair to suggest that Gettier missed the point of model-building, for the analyses of knowledge he was refuting were not intended as models; they were intended as statements of exceptionless necessary and sufficient conditions for knowledge, to which counterexamples were indeed apt. However, if epistemologists and other philosophers start aiming to build good models rather than provide exceptionless analyses, different forms of criticism become appropriate.

Models can also play a role in the criticism of would-be universal generalizations. If we are willing to dismiss theories on the basis of one-off negative verdicts in a single type of thought experiment, as with Gettier cases, we risk sometimes dismissing true theories because a glitch in the human cognitive system causes us to deliver mistaken verdicts in those thought experiments (Alexander and Weinberg 201X). A robust methodology should have ways of correcting such errors, even granted that thought experimentation is in general a legitimate method. After all, sense perception is a legitimate method for gaining knowledge, but we still need ways of catching and correcting perceptual errors. Elsewhere, I have argued that theoretical considerations about models of epistemic logic lead one to predict failures of the justified true belief analysis of knowledge, independently of thought experiments (Williamson 2013, 2015). When the methods of thought experimentation and model-building converge on the same conclusion, it has more robust support than when it relies on either method alone.

Another respect in which rigorous-minded philosophers may find the method of model-building alien is that selecting and interpreting models is an *art*—in science as well as in philosophy. It depends on good judgment, honed by experience. One must distinguish simplifications which abstract away inessential complications from those which abstract away crucial features of the phenomenon, and genuine insights from mere artefacts introduced for mathematical convenience. This raises the general issue of realism versus instrumentalism, familiar from the philosophy of science. Which aspects of a model tell us something about reality itself, and which are there only as instruments of the model-building process? We should not expect to settle all such issues in advance. Sometimes the successes of a model may indicate that what originally looked like a mere artefact should instead be regarded as a genuine insight. Although we can expect good model-builders to be reasonably articulate in explaining why they have selected one model rather than another and drawn one conclusion from it rather than another, there is no foreseeable prospect of reducing their skills and expertise to mechanical rules.

Some philosophers may continue to find the methodology of model-building mysterious, and resist. How can we learn from models that embody assumptions we know to be false? How exactly are we supposed to decide which false assumptions are legitimate? The short answer is: in the same way as the natural and social sciences. A full answer will be hard to articulate. Nevertheless, accumulating experience of model-building in philosophy provides good evidence that it does work.

Conclusion

Model-building already plays a significant role in contemporary philosophy. One neglected form of progress in philosophy over the past fifty years has been the development of better and better formal models of significant phenomena. It shares that form of progress with the natural and social sciences. Philosophy can do still better in the future by applying model-building methods more systematically and self-consciously. Although it is neither likely nor desirable for model-building to become the sole or even main philosophical method, its use enhances the power and reliability of philosophical thinking.

Note

I have presented earlier versions of this material at the universities of Athens, Cologne, Michigan (Ann Arbor), Olomouc, Oxford, Peking University, Seoul National University, and the Inter University Centre in Dubrovnik, the University of Olomouc; thanks to all the audiences for helpful questions and comments.

References

Alexander, Joshua, and Weinberg, Jonathan. 201X. "The "Unreliability" of Epistemic Intuitions." In *Current Controversies in Experimental Philosophy*, edited by Edouard Machery and Elizabeth O'Neill. London: Routledge.

Arntzenius, Frank. 2008: "Gunk, topology, and measure," *Oxford Studies in Metaphysics*, 4: 225-247.

Carnap, Rudolf. 1947: *Meaning and Necessity: A Study in Semantics and Modal Logic*. Chicago: Chicago University Press.

Ditmarsch, Hans van; Halpern, Joseph; Hoek, Wiebe van der; Kooi, Barteld. 2015: *Handbook of Epistemic Logic*. London: College Publications.

Gettier, Edmund. 1963: "Is justified true belief knowledge?" *Analysis*, 23: 121-123.

Godfrey-Smith, Peter. 2006a: "The strategy of model-based science." *Biology and Philosophy*, 21: 725-740.

Godfrey-Smith, Peter. 2006b: "Theories and models in metaphysics." *Harvard Review of Philosophy*, 14: 4-19.

Godfrey-Smith, Peter. 2012: "Metaphysics and the philosophical imagination." *Philosophical Studies*, 160: 97-113.

Hintikka, Jaakko. 1962. *Knowledge and Belief*. Ithaca, N.Y.: Cornell University Press.

Howson, Colin, and Urbach, Peter. 1993: *Scientific Reasoning: The Bayesian Approach*, 2nd ed. Chicago: Open Court.

Paul, Laurie. 2012: "Metaphysics as modelling: the handmaiden's tale." *Philosophical Studies*, 160: 1-29.

Russell, Jeffrey Sanford. 2008: "The structure of gunk: adventures in the ontology of space." *Oxford Studies in Metaphysics*, 4: 248-274.

Weisberg, Michael. 2007: "Who is a modeler?" *British Journal for the Philosophy of Science*, 58, 207-233.

Weisberg, Michael. 2013: *Simulation and Similarity: Using Models to Understand the World*. Oxford: Oxford University Press.

Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

Williamson, Timothy. 2013: "Gettier cases in epistemic logic." *Inquiry*, 56: 1-14.

Williamson, Timothy. 2014: "Very improbable knowing." *Erkenntnis*, 79: 971-999.

Williamson, Timothy. 2015: "A note on Gettier cases in epistemic logic." *Philosophical Studies*, 172: 129-140.

Wilson, Mark 2008: "Beware of the blob: cautions for would-be metaphysicians." *Oxford Studies in Metaphysics*, 4: 275-320.