MODELS AND REALITY

S YABLO

Contents

The title comes from a well-known paper of Putnam's (Putnam [1980]). The content is very different. Putnam uses model theory[1] to cast doubt on our ability to engage semantically with an objective world. The role of mathematics for him is to *prove* this pessimistic conclusion. I on the other hand am wondering how models can *help* us to engage semantically with the objective world. Mathematics functions for me as an analogy. Numbers among their many other accomplishments boost the language's expressive power; they give us access to recondite physical facts. Models, among their many other accomplishments, do the same thing; they give us access to recondite physical facts. This anyway is the analogy I will try to develop in this paper.

1. Applications

Mathematics is useful in physics. Frege was impressed by this: "It is applicability alone that raises arithmetic from the rank of a game to that of a science."

[1]The Lowenheim-Skolem Theorem.

Wigner found it mysterious, which is why he speaks of "the *unreasonable* effectiveness of mathematics" in physics. The mystery has only deepened with the attention in recent years to the *ways* in which math can be effective. Why should objects causally disconnected from the physical be so helpful in representing physical phenomena and making physical theories tractable? Why should math be such a good source of physical hypotheses? Why should it shed light on physical outcomes?

Before digging into these questions, consider *models* of the type appealed to in the natural sciences. They too are helpful in representing physical phenomena. They too make complex theories tractable. They too suggest hypotheses, and are apt to be cited in explanations. Why is there not a problem of the unreasonable effectiveness of models, as there is for mathematical objects?

The simplest answer is that we are dealing with a selection effect. Of all the technically eligible models that *could* be invoked, we focus, naturally, on the useful ones.

But, if that solves the problem for models, why does it not solve the problem for mathematical objects? The same selection effects are arguably at work with them.[2] Numbers are important because of their relation to counting and cardinality. Geometry grew out of land measurement problems, as the name suggests. Real numbers owe at least some of their prominence to being "complete" in the way space and time are thought to be complete. Calculus came to the fore in connection with Newtonian mechanics. Cantor's theory of the infinite grew out of calculus problems to do with integrability. Add to this that scientific models are causally independent, in most cases, of the phenomena that they model, and the contrast is hard to make out. The utility of models begins to seem similarly puzzling to that of mathematical objects.

2. Isms

The dialectic is not so different either. One popular theory of mathematical applications is *instrumentalism*. Numbers are useful, according to instrumentalists like Field, not for what they let us *say*, but what they let us *do*—shorten proofs as it might be:

> even someone who doesn't believe in mathematical objects is free to use mathematical existence-assertions in a limited context: he can use them freely in deducing nominalistically-stated consequences from nominalistically-stated premises (Field [1980], 14).

What about models? They too can be used in a purely instrumental way. There are

> 'probing models', 'developmental models', 'study models', 'toy models', [and] 'heuristic models'. The purpose of such model-systems is

[2] Balaguer [1998], Pérez Carballo [2014]

> not to represent anything in nature; instead they are used to test and study theoretical tools that are later used to build representational models. (Frigg [2010b])

Another leading approach to applications is *structuralism*. Structuralism about applications has been advocated in the philosophy of mathematics by Shapiro (Shapiro [1983]) and in the theory of modeling by van Fraassen and others (Van Fraassen [2006], Andreas and Zenker [2014]).

Fictionalism is an old standby where mathematical applications are concerned (Papineau [1988], Balaguer [1996]), Leng [2010]). It has recently made the jump to models:

> A natural first description of [frictionless planes, e.g.] is as fictions... They do not exist, but at least many of them might have existed, and if they had, they would have been concrete, physical things, located in space and time and engaging in causal relations. Though imaginary, these things are often the common property of a community of scientists. They can be investigated collaboratively. Surprising properties might be uncovered by one investigator after being denied by another. In their status, though not their role, they seem analogous to the fictions of literature. ((Godfrey-Smith [2006])

Fictionalism has morphed in recent years into *figuralism*, which sees numbers as creatures of metaphor, or (in Walton's version) of prop-oriented make-believe (Yablo [2002], Yablo [2005]). The make-believe approach has been tried for models too. For Frigg (description taken from Levy [2015]),

> Models are Waltonian games of make-believe. A set of equations or a mechanism sketch is a prop that, together with the rules relevant for the scientific context, determines what those engaging with the model — the game's participants — ought to imagine...[The] text and equations aren't, in this view, a description of an imaginary entity but a prescription to imagine a ring-shaped embryo with the specified chemical makeup. Thus, there is no object... to which [the] equations somehow correspond. There are only inscriptions on a page which function as instructions for the imaginations of modelers.

Levy too thinks of models as rules for the imagination. But the props in his view "are the real-world target phenomena" we are trying to understand. The role of the model is to

> portray a target as simpler (or just different) than it actually is. The goal of this special mode of description is to facilitate reasoning about the target. In this picture, modeling doesn't involve an

> appeal to an imaginary concrete entity, over and above the target. All we have are targets, imaginatively described (Levy [2015])

Of course, different applications may call for different approaches; we may want to be instrumentalists here and fictionalists there. More on this in a moment. Let's try to bear the above analogies in mind as we turn from theories of applicability to models as entities in their own right.

3. Types of Model

What is a model? If you think that this is not the most important question to be asking, you are probably right. But we need to say *something* about it, for the word is genuinely ambiguous. Model *citizens* are paragons or exemplars of good citizenship. *Role* models are figures worthy of emulation. *Fashion* models are, well, you know. A certain Joseph Bell was reportedly the model for Sherlock Holmes. *Car* models are things like the Ford Cortina and Fiat Panda.

So far, so irrelevant to scientific modeling. Model *cars* are a bit more like it. These stand in for real cars, and serve as a guide—in wind tunnel experiments, for instance— to real cars' properties. Likewise the wind-up models of the solar system encountered in science class; they stand in for, and are a guide to the properties of, the actual solar system. Model solar systems and the like are valued for the light they shed on whatever it is that they model.

Models serving as a guide to real properties are called *representational*; they will be our main focus. Not all models are representational, as already noted. Some may be for playing around with, to get a *feel* for certain real systems. Some may be valued for the hypotheses they suggest. Some may play a proof-of-concept role. Morrison and Morgan list some further possibilities:

> Just as we use tools as instruments to build things, we use models as instruments to build theory (Morrison and Morgan [1999], 18) Models are often used as instruments for exploring or experimenting on a theory that is already in place (ibid, 19) [M]odels are instruments that can both structure and display measuring practices (ibid, 21) The [class] of models as instruments includes those that are used for design and the production of various technologies (ibid, 23)

Given all these alternatives, why the focus on *representational* models? First because they're central to the scientific project; science aims, so it is said, at the accurate representation of real systems. Second because there is work in the philosophy of mathematics we'd like to draw on, which construes numbers, et al, in representational terms.

Representational items go hand in hand, however, with things represented. These oddly enough are apt to be called models too; think for instance of artist's models, or the solar system as the model for the gadget in science class.[3] So although representational models are the focus, room will also have to be made for models of the thing-represented sort; the key question about the former is, after all, how they relate to the latter. Models of the thing-represented sort are sometimes called *targets* or *target systems* since they are what representational models are aimed or directed at. Models sans phrase will be representational unless otherwise indicated.

4. Types of Truth

Representational models, like models in general, come in lots of varieties. There are *scale* models, like the wind-up solar system or the balsa wood wing in the wind tunnel. These are actual concrete particulars. The Bohr model of the atom is a *type* of concrete particular, a type that is not actually instantiated. Ideal models, frictionless planes and the like, are *would-be* concrete particulars; that is what they would be if they existed (Godfrey-Smith [2009]). The National Weather Service's climate models are computer simulations. Models of computing, like Turing machines or pushdown automata, might be seen as abstract particulars, or types of abstract particular. The model of electric current as water flow is an analogy. The Lotka-Volterra model of predator-prey relations is a set of equations.[4]

If there is anything tying these various types of model together, it is not their ontological category. That being said, we can without *too* much violence force most of them into the hypothetical-concrete-particular mold. The role of the Locke-Volterra model is played by concrete populations described by the equations. The models associated with a computer simulation of El Nino are the concrete meteorological processes that satisfy the simulation's assumptions. This puts models in many cases into the same metaphysical category as target systems, which doesn't matter now but will come in handy later.

A better bet for the common element would be how they *function*— what they do for us. Once again, models in general have *lots* of functions. They are used for testing and prediction, as aids to calculation and visualization, to manage complexity and facilitate understanding. One can say more about the function of *representational* models (this may be stipulative). These are meant to

(1) improve our access to the reality being modeled—the target system,
(2) by providing an epistemically accessible substitute,
(3) information about which translates into information about the target.

[3] I do not include Joseph Bell here because Holmes is under no obligation to be true to Joseph Bell, nor is he used as a guide to Joseph Bell's properties.

[4] The models employed in philosophical logic, like Kripke's fixed point model of a semantically closed language or the Bayesian model of belief update, are constructions or construction techniques; this may apply to philosophical models more generally (Godfrey-Smith [2006], Paul [2012], Williamson [2016]).

Models aspire, if the characterization is right, to be somehow a reliable guide to—I will say true to—the facts. One of the fundamental issues about models is to see what "true to the facts" could possibly mean here. Balsa-wood wings and computer simulations are not even *apt* for truth, it would seem, for they don't *say* anything. Truth is a property of statements or claims, not pieces of wood or programs. But let us push a little further.

The property reserved to statements is *declarative* truth, the kind Aristotle and Tarski talked about. Declarative truth is on some views not the only kind of truth out there. If we ask for a true copy of some document, or call a portrait true to its subject, we seem to be talking about accuracy or lifelikeness or fidelity. How far these should be considered kinds of *truth* is open to doubt. (Rooms with more and better portraits in them do not seem to contain more *truths*.) That is not important for our purposes. It's enough for us that (i) *declarative* truth is a kind of truth, and (ii) declarative truths are at least *part* of what we hope to gain from our (representational) models. The question either way is, how will this be possible, if models are not candidates for declarative truth?

Put just like that, this question is not very difficult. You might equally ask how we hope to learn truths from *newspapers*, or by *inference*, or by asking for *directions*; these things are not apt for declarative truth either. Why should they be? It's enough if there are semantic truths in the neighborhood, to which newspapers (inferences, directions) provide access. And indeed there are: the editor's claim on *behalf* of a newspaper that it is largely accurate, or your informant's claim on behalf of the directions she gives that they will get you to your destination. With models too, there is a candidate for semantic truth in the neighborhood: the theorist's claim on behalf of a model that it is faithful in such and respects to its target.

All of this is roughly in line with RIG Hughes' DDI theory of how models function (Hughes [1997]), characterized here by Frigg:

> According to [the] DDI account of modeling, learning takes place in three stages: denotation, demonstration, and interpretation. One begins by establishing a representation relation (denotation) between the model and the target. Then one investigates the features of the model in order to demonstrate certain theoretical claims about its internal constitution or mechanism; i.e., one learns about the model (demonstration). Finally, these findings have to be converted into claims about the target system; Hughes refers to this step as 'interpretation.'[5] (Frigg and Hartmann [2005], 744ff)

[5]Hughes is not aiming here for an analysis: "I am not arguing that denotation, demonstration, and interpretation constitute a set of speech acts individually necessary and jointly sufficient for an act of theoretical representation to take place. I am making the more modest suggestion that, if we examine a theoretical model with these three activities in mind, we shall achieve some insight into the kind of representation that it provides (Hughes [1997] S329)"

Our concern is mainly with the interpretation stage: converting findings, or more generally claims, about the model into claims about the target system—more specifically, with the ways in which claims can be "about the target system."

5. Targeting (I): The Role of Math in Real Content

The model matches in such and such respects the target is a semantic truth, all right. But I wonder if it is the kind of semantic truth we were after. To say that the target system resembles the model is to speak in part of the model. And we wanted a claim about the target, or more broadly the world.

You may say, why should it not be about both? I will give some reasons in a moment, but the problem in a nutshell is that although inquiry *avails* itself of models, it should not be (in cases of interest) *about* models. There should be the possibility, at least, of wringing truths entirely about the target out of properties of the model.

This is nothing special about models. Self-directedness is unfortunate with lots of representational devices. Take graphs, or barometers. Inquiry avails itself of them, but does it aim at truths *about* graphs and barometer? These would be truths like, *The graphs in Feynman's Lectures are largely accurate*, or *Air pressure as measured by barometers falls in a thunderstorm.* Surely not. One is hoping ultimately to be *using* representers as a means of access to information about the world.

This is admittedly just an appeal to intuition. But we can do better, for a similar issue comes up in the literature on mathematical applications. The assumption there is that math-involving talk is in a sense *hyperbolic*. One *quasi*-asserts an S directed in part at numbers, in order to *really* assert a weaker claim $\rho(S)$ that is not about numbers at all. I might quasi-assert that the number of cells in this petri dish is doubling every day in order to really assert that there are always twice as many cells as the day before—which is, to anticipate a little, the part of S about concreta. What S says about concreta is S's *real content* in a setting where we are talking about the physical.

Why should the real content have to be wholly about concreta? The reason is that representational devices, including numbers but not only them, are "out of place" in certain contexts. The real content has to be number-free to have, in certain contexts, the right truth-conditional effects. Numbers are out of place because allowing them into the real content winds up falsifying a larger claim that ought to come out true.[6]

For a sense of how this might work, consider the context of *causal explanation*. Field's case for nominalism in *Science Without Numbers* relies on the idea that explanations ought to be "intrinsic":

[6]Arguments of this type are developed at greater length in Yablo [2001, 2002].

> If we need to invoke some real numbers like 6.67 x 10^{11} (the gravitational constant in [SI units]) in our explanation of why the moon follows the path that it does, it isn't because we think that that real number plays a role as a cause of the moon's moving that way...The role it plays is as an entity extrinsic to the process to be explained, an entity related to the process to be explained only by a function (a rather arbitrarily chosen function at that) (Field [1980], 43)

The *real* reason the moon follows that path is to do with the strength of the forces acting on it, not their numerical representation.

Suppose Field is right that explanations should confine themselves to the entities actually doing the work. Allusions to numbers could then make an explanation defective—not to split hairs, let's just say *false*—in roughly the way that the allusion to God casts doubt on *The patient recovered because God knows she was given antibiotics*. Numbers are unwelcome in these contexts because they are extrinsic to the causal scene. The bag ripped because it had too many apples in it, not because a certain number (the number of apples in it) was too large. It is not that numerals can never *appear* in X *because* Y, or that they can never *influence* the real content of X or of Y.[7] The suggestion is only that numbers should not *participate* in the real contents, if this would violate some plausible version of the intrinsicness constraint.

If numbers are indeed objectionable in causal/explanatory contexts, perhaps the real content should treat them as existing only *according to a certain story*, the story of standard math. But, the story of standard math is just as extrinsic to the scene as the numbers of which it treats. The bag didn't rip because a certain number was too large *according to standard math*, any more than it ripped because a certain number was too large.[8]

Or consider *nomological* contexts. According to Galileo's Law of Falling Bodies, $d(t)$ — the distance a dropped object falls in t seconds — is proportional to t^2. Suppose we are convinced for broadly Fieldian reasons that the real content of $d(t) \propto t^2$ in this setting does not involve numbers or numerical operations (like squaring); the law treats of concrete objects, not mathematical ones. Matters are not improved by putting the numbers under a story prefix, for natural laws know nothing of stories either. If Galileo's Law is really to be a law, the real content of $d(t) \propto t^2$ should not involve the story of standard math.

Or contexts of *understanding*. I may need to know some math to understand Galileo's Law, in its standard formulations. I do not, however, need to know what standard math *is*, to understand it. Since I *do* apparently need to know

[7] *Members of Congress cannot be paired off one-one because the number of them is odd.*

[8] I am using "the story of standard math" loosely to allow the importation of truths about non-mathematical objects. Incorporation of real or apparent truths into the content of a fiction is standard operating procedure. See Walton on the *Reality Principle* and the *Mutual Belief Principle* (Walton [1990],144ff).

which math is standard to understand, *According to standard math, distance fallen is proportional to the square of the time elapsed*, the latter is not a very good candidate for the real content of Galileo's Law. *Agreement* contexts are similar. People agree, I take it, on Galileo's Law; at least there is no obvious obstacle to our agreeing. A potential obstacle emerges, though, if the Law has the story of standard math in its real content. For who is to say we are working with the same version of the story?

Modal contexts put a different kind of pressure on real content . The number of $(F \vee G)$s is *bound* to be even, if the number of Fs = the number of Gs, and the number of $(F \wedge G)$s = 0; it could not have been otherwise. But it could (perhaps) have been otherwise according to standard math, for standard math could have been different.

Consider finally *epistemic* contexts. It is supposedly a priori that if the number of Fs = the number of Gs, and the number of $(F \wedge G)$s = 0, then the number of $(F \vee G)$s is even. But do we know a priori that this is so according to standard math? No, because we know do not know the content of standard math a priori.[9] Again, we seem to know a priori that a set's subsets outnumber its members. But, this holds only on a combinatorial conception of set. And it is somewhat of a historical accident—it was not anyway inevitable—that set theory developed in that direction.

6. Targeting (2): The Role of Models in Real Content

These are some of the problems that arise if representational devices are written into the real content of math-infused statements. It would be surprising if similar problems did not sometimes arise when representational models are written into the real content of statements of model-based science.

Suppose we are working with a purely gravitational model of the solar system in which planets interact exclusively with the sun. (I will use α for actual systems and ω for models.) And suppose that the hypothesis about α that we access by quasi-asserting S—asserting it in reference to ω— is not entirely about α but involves also ω. It is, let's say, the hypothesis that ω is similar in respect R to α. To use our earlier terminology, $\omega R \alpha$ is the real content $\rho(S)$ of our quasi-assertion that S.

What kind of trouble is caused, in what contexts, by the real content's alluding not only to the target system α but also the model ω?

Start as before with causal/explanatory contexts. The "effect" is that planets speed up on approaching the sun. We'd like to explain it with Kepler's Second Law: A planet always sweeps out the same area in the same amount of time. This is not strictly true, though, of α. The real content $\rho(S)$ of Kepler's Law *is* true, however, and we look to it for the explanation. We look in vain if the real content

[9] To know what is true in the Holmes stories, one has to look at the stories.

is of the form $\omega R\alpha$, because the model does not participate in the Earth's reasons for speeding up when it approaches the sun; it figures at most in the representation of those reasons.[10]

A law's holding without exception in ω is meant to tell us something lawful about actuality. So, it reflects a robust fact about the solar system α that planetary orbits are elliptical in ω. But it reflects no deep fact about the solar system that it is R-related to ω. This for two reasons, one pertaining to R and one to ω. Why would an astronomical law of *this* world bring in a system that exists only in other worlds, or α's relations to this nonexistent structure?

And so on. If the real content is $\omega R\alpha$, then it is accessible only to those acquainted with ω. To agree on the real content of S, we have to be working with the same ω. If the real content is $\omega R\alpha$, information gleaned from multiple models does not paint a unified picture: the most we can say is that α resembles this model in one respect, that one in another, a third in a third respect, and so on.[11]

7. Actuality

The target system α is supposed in most cases to be real; it is part of the actual world. It simplifies mattersd to treat it as identical to the actual world, on the understanding that truths S about the model translate into truths (true real contents) $\rho(S)$ that pertain only to the bits of actuality that are being modeled. The model itself is presumed *not* to be real; it is part of a counterfactual world. I propose again to treat it as identical to that world, on the understanding that S speaks only to the bits that do the modeling. There are plenty of other options here. The target system could be a mini-world, for instance, or a situation, or a set of worlds with only the mini-world in common; and similarly for the model.

I am not going to fret too much about the ontology of models and target systems because the action is really elsewhere. Truths about ω, we said, are supposed to translate into truths $\rho(S)$ about α. As Frigg puts it,

[10] Bokulich [2011]

[11] See Weisberg [2007, 2012] for multiple models of the same target, and multiple targets for the same model (e.g. a compressed spring device governed by Hooke's Law stands in for harmonic oscillators generally). Admittedly a unified picture is not always desirable, or possible; one of the glories of model-based theorizing is supposed to be that it takes this in stride. Weisberg quotes Levins: "The multiplicity of models is imposed by the contradictory demands of a complex, heterogeneous nature and a mind that can only cope with few variables at a time; by the contradictory desiderata of generality, realism, and precision; by the need to understand and also to control; even by the opposing aesthetic standards which emphasize the stark simplicity and power of a general theorem as against the richness and the diversity of living nature. These conflicts are irreconcilable. Therefore, the alternative approaches even of contending schools are part of a larger mixed strategy. But the conflict is about method, not nature, for the individual models, while they are essential for understanding reality, should not be confused with that reality itself" (Levins [1966]).

> Models not only represent their target; they do so in a clearly specifiable and unambiguous way, and one that allows scientists to 'read off' features of the target from the model...we study a model and thereby discover features of the thing it stands for. We do this by first finding out what is true in the model-system itself, and then translating the findings into claims about the target itself. (Frigg [2010a])

But now, where are we to look for the translation manual? It should be possible, on Frigg's picture, to identify $\rho(S)$ on the basis of S, α, ω, and ω's relation to α, henceforth R. His schematic solution is that

> ω comes with a key K specifying how facts [S] about ω are to be translated into claims [$\rho(S)$] about α (Frigg [2010a])

This is not very informative, as Frigg says himself:

> [T]here is much more to be said ... than is contained in [the conditions given]—they are merely blanks to be filled in every particular instance. Thus, the claim that something is a representation amounts to an invitation to spell out how exactly ω comes to denote the target system α and what K is (Frigg [2010a]).

The key of course is none other than the sought after translation manual. How it might be found is to be determined on a case by case basis. Contessa tries to say something more general. The model-user

> adopts an interpretation of the [model] in terms of the target ...and this interpretation provides the user with a set of systematic rules to "translate" facts about the vehicle into (putative) facts about the target (Contessa [2011])

The interpretation sounds at first like a way of *seeing* the model that treats certain aspects as representational and others as adventitious. This is tantamount in Contessa's view to adopting a set of systematic rules that translate facts about the vehicle into facts about the target. But, seeing the model a certain way falls far short of learning systematic translation rules. If on the other hand the interpretation involves more than a way of seeing—if it is defined so as to provide rules— then it is not clear that actual model-users ever adopt interpretations, or how they find the ones that are worth adopting. Adam Morton pushes back further, to the reasons for our initial choice of model. A candidate model, to attract our attention, must (among other things)

> give predictions that are reliable in specific but rarely explicitly specified respects. (Morton [1993], 663)

And now one might argue as follows.[12] The model-chooser should have some idea, surely, of the type of reliability that motivated her choice. And how does knowing the relevant type of reliability fall short of knowing a translation manual mapping truths S about ω to truths $\rho(S)$ about α? Again we face a dilemma. Either a "type of reliability" is tantamount to a translation manual, or it falls short, providing, as it might be, a specification of the range over which the model is trustworthy, without a specification of the message worthy of trust. If it falls short of a translation manual, then the argument never gets off the ground. If it suffices for a translation manual, then while model-choosers may have "some idea" of how the model is reliable, and "some idea" of how to wring predictions out of it, they do not normally know what the manual *says*; the predictions licensed by a property of the model are, as Morton says, rarely explicitly specified.

To narrow the search space it might help to consider the *form* that reliable predictions or claims must take. A number of proposals have been made, or hinted at, about the implication for α of the fact that ω satisfies S. The implication might be that

(1) α has features analogous to the features that ω needs to satisfy S
(2) a model of S has a part ("appearances," e,g,) isomorphic to part of α[13]
(3) α is such as to make S true in a certain story, the "story of ω"[14]
(4) an analog for α of S answers an analog for α of the question S addresses

These, however, all allude in one way or another to ω, which we saw above to cause trouble. A $\rho(S)$ partly about ω is vulnerable to the objections raised at the end of section 2 to a real content that portrays α as resembling ω in a certain respect. Look for example at (1). No one can understand the features that (1)

[12] Not that Morton argues this way. I'm just mining his work for constraints on a translation manual.

[13] Van Fraassen [1980], 64.

[14] Frigg [2010b], Toon [2010], Godfrey-Smith [2009], Levy [2012] and Levy [2015] all explore the idea of treating truth in the model as a kind of fictional truth, or pretense-worthiness in a Walton-style prop-oriented make believe game (Walton [1993]). But although these authors cite Walton, they seem mostly—Levy [2015] is an exception—to ignore his picture of how make-believe can be used in the cause of real-world representation. The point of prop-oriented make-believe for Walton is to give information about the props—the real-world items determining what is to be pretended. To utter S in the context of a game $\mathcal{G}$ is supposed to be a way of representing the props as in a condition to make S pretense-worthy in that game. If we're trying to represent the solar system, the prop should be the solar system. Suppose we were to identify props with model descriptions, as Frigg appears to; this in Walton's scheme means that the point of uttering S is to give descriptions of model-descriptions. Likewise if the props are the models themselves. The point of uttering S in connection with a representational device is not to give information about the device; that would be like treating *Crotone is in the arch of the Italian boot* as a guide to the model of Italy whereby Italy is a boot. It's to give information about the world. (To be sure, one can follow Walton in his theory of make-believe entities, without following him on the representational point of such entities. As far as I know only Levy goes all the way.)

attributes to α if they are not acquainted with ω. (1)-type information is shareable only between theorists working with the same ω. Possession of features analogous to those by which ω satisfies S does not cause possession of features analogous to those by which ω satisfies S'. The model's role should be to induce a content in which it does not itself figure.

8. Translation

The problem as we've been conceiving it so far (but not much longer) is how to translate a truth S about the model into a truth $\rho(S)$ about the world, that is, a truth full stop. In schematic form:

N:
S holding in a model suitably related to α testifies,
not to the truth of S itself in α, but
the truth in α of a hypothesis $\rho(S)$ suitably related to S.

This way of setting the problem up requires us to identify $\rho(S)$, however—which has been proving difficult. We should ask ourselves whether a translation of S into α-ese is really necessary.

You might think it clearly *is* necessary, since S untranslated is (normally) false about α. But there are other alethic commendations we can give to a sentence besides truth. Perhaps all that S aspires to is to be *partly* true in α—true apart from an issue we're properly ignoring. Or perhaps it aspires only to be true about a certain *aspect* of α. Or it might aim to be true about α where a subject matter of particular interest is concerned. The idea more generally is that, rather than attempting to translate S into a claim that is wholly true, we might try to scale "wholly true" back to a compliment that S is worthy of as it is.

How would this work in practice? Let S be Kepler's Law that the planets trace elliptical orbits. This law, although false in α, is true in a model with gravitational forces only and a single planet revolving around a massive central body like the Sun.[15] Now, what is the fact about α that is indicated by Kepler's Law holding in the model? That Kepler's Law is true in α about the matter of `planetary motion due to centripetal gravitational forces`. Another fact indicated is that Kepler's Law is *roughly* true in α, or true about the matter of `planetary motion give or take a certain fudge factor`. Schema **N** thus gives way to

M:
S holding in a model suitably related to α testifies,
not to S's truth simpliciter in α, but
its truth in α as far as subject matter $\mathtt{m}$ is concerned

[15] See Cartwright [1983] for the idea that laws hold only in simplified models, and Yablo [2014], 84-5 for discussion.

Continuing along the same lines, gases are modeled as collections of randomly moving, perfectly elastic, point-sized, non-interacting particles trapped in a perfectly smooth and rigid container. Pressure rises in such a model with the number and speed of the particles colliding with the container's wall. This cannot be said of real gases, since gas particles are not literally point-sized, and they interact; the wall is not really smooth or rigid, it is made of atoms; the particles do not really collide with these atoms but are repelled by them electromagnetically.

So what is going on here? I don't honestly know but people tell me things like the following. The number and speed of collisions stand in for the kinetic energy of the impinging particles. The particles are moving randomly so that the mean kinetic energy of impinging particles is constant throughout. They are imagined as point-sized lest some of this energy be lost to rotation. The original statement may not be true *overall* in α, but it is true of an *aspect* of α, namely, `pressure as a function of mean translational kinetic energy`.

Consider next Schelling's grid model of housing preferences. This model has families relocating if and only if fully *three quarters* of their neighbors are of a different race; they are content in other words to have three same-race neighbors and five of a different race. The surprising thing is that racial segregation results after a few iterations of this enlightened-seeming process. The emergence of racially divided neighborhoods shows how segregation could have arisen "innocently," out of a desire not to be outnumbered 3:1 in the neighborhood (Schelling [2006]).

What is the lesson of Schelling's model for our world? Not that racial segregation results from dislike of racial isolation. This is true in the model but not necessarily in the world. The lesson is that racial segregation *could* result just from dislike of isolation, as far as the statistical evidence goes. Or, to say it a bit differently, the original statement gets a certain *aspect* of our world right, namely, `how much racial animus is required for segregation`. It tells us how *could* have arisen innocently, compatibly with the statistical data that seem to rule it out.

A final example is Fisher's equillibrium model of 1:1 sex ratios. Fisher claims, as I understand it, that there are forces at work that would lessen numerical disparities if the numbers got out of whack. These forces operate, in his model, by exerting selective pressure on a hypothetical gene that biases the sex of offspring toward males.

Suppose that females were in the majority at some point in evolutionary history. Newborn boys would then have better mating prospects than newborn girls, so that those with the male-favoring gene will have on average more grandchildren. The male-biased grandparent may not have more children, but there more boys among their children. These on account of their rarity will mate more frequently, putting copies of the male-tending gene into more grandchildren. And so on and so on, for as long as males have the mating advantages conferred by being in the minority. (Mutatis mutandis if males were more common.)

Fisher's story does not have to be true in all details (or at all) to shed light on actual sex ratios. It serves at the very least a as a proof of concept for the idea of selectional pressures favoring genes that deskew unequal populations. But it could conceivably be more than that, for we might get similar dynamics if sex ratios get out of whack for other reasons.

Suppose that some resource (food, say) is divided equally between the sexes. Then as the male population dwindles, each man becomes better fed, which increases longevity and thus the number of males. Or perhaps men protect women from predators, while women protect men from poisonous fruit. As males become more uncommon, women are increasingly preyed upon, bringing their portion of the population down to male levels. As females become more uncommon, men increasingly die of poisoning, bringing their representation down to female levels.

These are just-so stories, of course. But there would if any such story worked be a compliment we could pay to Fisher's theory *even if it was false*: the theory is true about `the tendency of unequal sex ratios to correct themselves.`

The structure of all these cases seems broadly similar. We have a statement S that is true in ω but false in α. S's truth in ω signals somehow its truth in α about a subject matter `m` that ω and α agree on. Altogether then: for S to be in ω indicate its truth in α about—and here we fill in the appropriate subject matter:

(1) *The Earth traces an elliptical orbit*
is true in ω and false in α; but true in α about
`planetary motion due to central gravitational forces`.

(2) *Pressure rises as point particles collide harder with the wall*
is true in ω and false in α; but true in α about
`how mean translational kinetic energy relates to pressure`.

(3) *Desire for >2 same-race neighbors in 8 leads to 7 such neighbors in 8.*
is true in ω and false in α; but true in α about
`how fear of racial isolation can lead to racial segregation`

(4) *Selective pressures on sex-bias genes make uneven ratios unstable*
is true in ω and false in α; but true in α about
`the instability and self-correctingness of uneven sex ratios`

9. Subject-Matter

The role of subject matter for us is to be the kind of thing sentences can be true about in a world, even if they're not true outright in a world. It doesn't matter for this purpose whether subject matters are "of sentences," or whether there is such a thing as *the* subject matter of A' for particular sentences A.[16] Subject matters for our purposes can be entities in their own right. David Lewis initiated

[16] I happen to think there is such a thing, but that's another story (Yablo [2014])

this approach in "Statements partly about observation" (Lewis [1988]). `The 19th century`, for instance, is a kind of subject matter, for Lewis. It's the kind he calls *parts-based.*

A subject matter `m` is parts-based, if for worlds to be alike with respect to `m` is for corresponding parts of those worlds to be intrinsically indiscernible. `The 19th century` is parts-based because worlds are alike with respect to it if and only if the one's 19th century is an intrinsic duplicate of the other's 19th century. The 19th century, note, is not to be confused with `the 19th century`. The first is a part of one particular world (ours), or of its history. The second is a way of *grouping* worlds according to what goes on in their respective 19th centuries.

This approach is not sufficiently general, Lewis observes. Take the matter of how many stars there are. There is no "star-counter" part of the universe, such that worlds agree in how many stars they contain if and only if the one's counter is an intrinsic duplicate of the other's. Facts about how many stars there are are not stored up in particular spatio-temporal regions.

Or consider the subject matter of `observables`,[17] which van Fraassen uses to define an empirically adequate theory. This is prima facie a parts-based subject matter, like `the Sun`. Worlds are observationally equivalent just if their observables—whatever in them can be seen, or heard, or etc—are intrinsically alike. But again, dirt can be seen, and among dirt's intrinsic properties are some that are highly theoretical, for instance, the property of being full of quarks. It is not supposed to count against a theory's empirical adequacy that it gets subatomic structure wrong.

`Observables`—what an empirically adequate theory should get right—is best regarded as a *non*-parts-based subject matter, like `the number of stars`. Worlds are alike with respect to `observables` if they're observationally indistinguishable; they look and feel and sound (etc) the same.[18]

10. Partitions

A parts-based subject matter, whatever else it does, induces an equivalence relation on, or partition of, "logical space."[19] Worlds are equivalent, or cell-mates, if corresponding parts are intrinsically alike.

[17] Lewis calls it `observation.`

[18] What becomes then of the idea, seemingly essential to constructive empiricism, that T need only be true to the observable part of reality, if `observables` does not correspond to a part of reality? See Chapter 1 of Yablo [2014].

[19] Lewis [1988]. An equivalence relation $\equiv$ is a binary R that's reflexive (everything bears R to itself), symmetric (if x bears R to y, then y bears R to x), and transitive (if x bears R to y and y bears it to z, then x bears R to z). A partition is a decomposition of some set into mutually disjoint subsets, called cells. Equivalence relations are interdefinable with partitions as follows: x's cell $[x]$ is the set of ys equivalent to x; $x{\equiv}y$ if they lie in the same cell.

A *non*-parts-based subject matter, however, *also* induces an equivalence relation on logical space: worlds are equivalent, or cell-mates, just in case they are indiscernible where that subject matter is concerned. If `m` is `the number of stars`, $\equiv_{\mathtt{m}}$ is the relation one world bears to another just if they have equally many stars. But then, if one wants a notion of subject matter that works for both cases, let them be not *parts* but *partitions*. The second notion subsumes the first while exceeding it in generality.

So, to review—one starts out thinking of subject matters as parts of the world, like the western hemisphere or Queen Victoria or the 19th century. These then give way to world-*partitions*, which are ways of grouping worlds. Should the grouping be on the basis of goings on in corresponding world-parts, we get the kind of subject matter that, although still thoroughly partitional, looks back to world-parts for its identity-conditions.

A subject matter (or *topic*, or *matter*, or *issue*) on this view is a system of differences, a pattern of cross-world variation.[20] Where the identity of a set is given by its members, the identity of a subject matter is given by how things are liable to change where it is concerned:

SM:
$\mathtt{m}_1 = \mathtt{m}_2$ iff worlds differing on the one differ also with respect to the other.

This might seem too abstract and structural. To know what $\mathtt{m}_1$ is as opposed to $\mathtt{m}_3$ doesn't seem to tell us what goes into a world's $\mathtt{m}_1$-condition, as opposed to its $\mathtt{m}_3$-condition. Surely though I do grasp a subject matter `m`, if I can never tell you what the proposition $\mathtt{m}(w)$ that specifies how matters stand in w where `m` is concerned.

But, subject matters as just explained *do* tell us what w is like where `m` is concerned. The proposition we're looking for is meant to be true in all and only worlds in the same `m`-condition as w; on an intensional view of propositions, it *is* the set of worlds in the same `m`-condition as w. That proposition is already in our possession. To be in the same `m`-condition as w is to be `m`-equivalent to w, and the set of worlds `m`-equivalent to w is just w's cell in the partition. A worlds `m`-cell is thus the proposition saying how matters stand in it `m`-wise.

Lewis writes `nos` for `the number of stars`. How do we find the proposition specifying how matters stand in a world where `nos` is concerned? Well, w has a certain number of stars, let's say a billion. Its `nos`-cell is the set of worlds with exactly as many stars as w. The worlds with exactly as many stars as w are th4us the ones with a billion stars. The worlds with a billion stars comprise the proposition that there are a billion stars. That it contains one billion stars sums

[20] Linguists have their own notion of topic; a sentence's topic/focus structure is something like its subject/predicate structure. Topics in the linguist's sense may or may not be reflected in a sentence's subject matter.

up w's `nos`-condition quite nicely. By transitivity of identity, its `nos`-cell sums up its `nos`-condition quite nicely.

11. Ways to Be

Where does this leave us? The subject matters in (**10**) seemed too abstract and structural to tell us what is going on `m`-wise in a given world. But each `m` determines a function `m`(...) that encodes precisely that information. It works backwards, too; one can recover the equivalence relation from the function, by counting worlds `m`-equivalent if they are mapped to the same proposition.[21] `m` can thus be conceived as (i) an equivalence relation—that's what it is "officially"—or (ii) a partitition, or (iii) a specification for each world of what is going on there `m`-wise. `The number of stars`, for instance, can be construed as a function taking each k-star world w to the proposition `There are exactly` k `stars`.

The problem may seem to recur at a deeper level. How are we to get an intuitive handle on the function `m`(...) taking worlds to their `m`-conditions? It's one thing if `m`(...) is introduced in the first place as specifying how many stars a world contains. But all we know of specification functions considered in themselves is that they are mathematical objects (sets, or partial sets, presumably) built in such and such ways out of worlds. It is not clear how we are to think about sets like this, other than by laying out the membership tree and describing the worlds at terminal nodes as best we can.

Each specification function `m`(...) has associated with it a set of propositions, expressing between them the various ways matters can stand where `m` is concerned. (A proposition goes into the set if it is `m`(w) for some world w.) The operation is again reversible: to find `m`(w), look for the proposition to which w belongs.

A subject matter can also be conceived, then, as (iv) a set of propositions. Sets of this type function in semantics as what is expressed by sentences in the interrogative mood. *Questions*, as they are called, stand to interrogative mood sentences Q as propositions stand to sentences S in the indicative mood.[22] To find a Q expressing a particular set of propositions, look for one to which those propositions are the possible answers. This Q gives us an immediately comprehensible designator for the set of propositions at issue.[23]

What, for instance, is the Q to which `There are exactly` k `stars`, for specific values of k, are the possible answers? It is *How many stars are there?* We are dealing, then, with the issue or matter of `how many stars there are`. What

[21] This won't work, of course, with any old function from worlds to sets of them. The proposition associated with w must be true in it; the propositions associated with different worlds should be identical or incompatible.

[22] I will sometimes use "question" sloppily as standing also for the sentences.

[23] By pointing us to the corresponding indirect question. The indirect question corresponding to *Do cats paint?* is *whether cats paint*. The indirect question corresponding to *Why do they paint?* is *why cats paint*.

is the question addressed by `She said BLAH to Francine`, for specific values of BLAH? It is *She said WHAT to Francine?* Thinking how to answer *What did she say to Francine?* is considering the matter of `what she said to Francine`. The question addressed by *Cats paint* is *Do cats paint?*; the corresponding subject matter is `whether cats paint`. The question addressed by *Cats paint to relieve tension*is *Why do cats paint?*. Pondering that question is reviewing the matter of `why cats paint`.[24]

12. Targeting (3): Directed Truth

That S holds in a certain model ω tells us, or can tell us, that S is true, period, where a certain subject matter is concerned. How does it do this? Recall that we earlier reconceived the target system as a world α, and decided to think of the model as a world or mini-world ω. This makes both into the kinds of thing that can be alike, or different, where a subject matter is concerned. And now we can elaborate schema **M** in two different ways:

M∃: S is true in α where `m` is concerned iff it is true simpliciter in a world ω that is equivalent to α where `m` is concerned.

For S to be true about `m` in α means (according to **M∃**) that S, should it be false in α, is at any rate not false because of how matters stand with respect to `m`. This admits of a simple test: S is not false about how matters stand `m`-wise iff one can make it true without changing how matters stand `m`-wise. The role, anyway one role, a model ω can play is to witness this possibility—the possibility of morphing our world into an S-world without disrupting the state of things `m`-wise.

Now, what kind of compliment we are paying S, when we call it true about `m`? Does truth about the subject matter under discussion make S "as good as true" for discussion purposes? Does "true about `m`" function in descriptions of @ the way truth simpliciter does? One has to be careful here.

Truth about `m`, considered as a modality, is possibility-like: A is true about `m` in a world just if it *could* be true, for all that that world's `m`-condition has to say about it. The logic of directed truth, on this view, can to some extent be read off the logic of possibility. This is fine for *some* purposes. Sometimes all we want from truth about `m` in α is that S *could* be true in the same `m`-conditions as obtain in α. This is how it works, for instance, with models of the solar system that replace planets with point masses stationed at each planet's center of gravity, or (as far as I can understand this issue) the models in solid state physics that have "quasi-particles" like phonons standing in for diffuse and large-scale vibrations.[25] The possibility-like notion of truth about `m` seems to suffice when, roughly speaking, there is only one ω for a given α, or S is satisfied by all the models of interest if any.

---

[24] See Silver and Busch [2006].

[25] See Gelfert [2003] and Falkenburg [2015]. Thanks to Jay Hodges for the example.

A particularly vivid example is invoking mathematical objects to increase expressive power. Take *The rabbit population in Australia was 27×2^n on the n^{th} day of 1866.*. This is not true outright in α, if α lacks numbers, as maintained by most if not all Australians. But it would still be true about how matters stood physically, if the rabbit population allowed this. That it could also be *false* under such conditions is not a problem. There is essentially only one way of fitting a physical world out with numbers.

Sometimes though it *is* a problem. A hypothesis and its negation can be possible at the same time. Similarly there is nothing to stop them from both being true about `m`. That a world's `m`-condition permits *each* of A and $\neg A$ to be true doesn't mean it permits them both to be true together. Truth about `m` is not agglomerative. Call this the phenomenon of *quasi-contradiction*.

How much of a problem is it if truths where `m` is concerned contradict each other? That depends on whether the contradiction is off-topic. (There is a problem only if A and $\neg A$ say contradictory things *about* `m`.) Take the statement that *The author of Principia Mathematica taught at Harvard*. This gets something right, in that Whitehead taught there. Its negation, however, also gets something right, since Russell did not teach at Harvard. There is nothing contradictory about only Whitehead teaching at Harvard!

That A and $\neg A$ can both be true about `m`, as long as they are consistent in what they *say* about `m`, is a nice outcome. But we pay a heavy price for it: truth about `m` is not closed under conjunction.[26] To obtain a notion of truth-about that *is* closed under conjunction, we need to put a universal quantifier in for the existential in **M∃**. This can done along Kratzerian lines (Kratzer [1977]) as follows.

> **M∀:** S is true in α where `m` is concerned iff it holds in all the best ωs that are equivalent to α where `m` is concerned.

Here I am imagining subject matters fitted out with an extra relation $>$; one of two `m`-equivalent worlds is better than another just if it better illuminates what is going on `m`-wise, by e.g., containing fewer distorting influences or irrelevant complications. The best `motion due to gravity`-equivalents of α, for instance, will have gravity as their *only* force and gravitationally induced motion as their only motion. This is important lest *Objects never move* comes out true in *our* world about `motion due to gravity`, thanks to a Zeno world where gravitational forces are exactly cancelled out by other forces.

13. Conclusion

Not a lot has been accomplished in this paper. A rough analogy has been developed between how numbers boost (can boost) expressive power and how models do. Sometimes the best way to get our point across is to advance a sentence S not as true full stop, but true about certain subject matter. we can do with

[26] Dorr [2010].

a false sentence. S's truth here in α about a subject matter $\mathtt{m}$ is analyzed as its truth simpliciter in a world ω that is just like ours where $\mathtt{m}$ is concerned. The principal difference is that ω in the numbers case is (assuming nominalism) *more* complicated than our world; it contains both a duplicate of α and a bunch of mathematicalia that are missing from α. Whereas ω in the case of models is *simpler* than our world, that being the whole point of working with models.

The main ideas about models have been that (i) translating truths in the model into truths about the target system is difficult, but (ii) there's an alternative: rather than trying to morph S into a truth about α, we can morph "true" into "true as far as $\mathtt{m}$ goes."[27] This provides a format but does not otherwise get us very far. The question now is, where do we look for a n $\mathtt{m}$ of which it's illuminating to know that S is true about it in $\alpha =$ the actual world? I don't know.[28]

References

Holger Andreas and Frank Zenker. Basic concepts of structuralism. *Erkenntnis*, 79(S8):1367–1372, 2014.

Mark Balaguer. A fictionalist account of the indispensable applications of mathematics. *Philosophical Studies*, 83(3):291–314, S 96 1996.

Mark Balaguer. *Platonism and Anti-Platonism in Mathematics*. Oxford University Press, 1998.

Alisa Bokulich. How scientific models can explain. *Synthese*, 180(1):33–45, 2011.

Nancy Cartwright. How the laws of physics lie. Clarendon, 1983.

Gabriele Contessa. Scientific models and representation. *The continuum companion to the philosophy of science*, page 120, 2011.

Cian Dorr. Of numbers and electrons. In *Proceedings of the Aristotelian Society*, volume 110, pages 133–181, 2010.

Brigitte Falkenburg. How do quasi-particles exist? In *Why More Is Different*, pages 227–250. Springer, 2015.

Hartry H. Field. *Science without Numbers: a Defence of Nominalism*. Princeton University Press, 1980.

Roman Frigg. Fiction and scientific representation. In *Beyond Mimesis and Convention*, pages 97–138. Springer, 2010a.

Roman Frigg. Models and fiction. *Synthese*, 172(2):251–268, 2010b.

---

[27]If you like, and I do, what S says about $\mathtt{m}$ is true.

[28]Let Q be the question about ω that is addressed by S, conceived as a partition of logical space. We are given that ω belongs to the S-cell of Q. We cast about for a question Q' about α one of whose cells (the one with α in it) overlaps the S-cell of Q. If one is found, this tells us that α resembles the S-worlds in the area of overlap in how it answers Q'.

Now let $\mathtt{m}'$ be the subject matter corresponding to Q'—the one such that the states of things $\mathtt{m}'$-wise are all and only the answers to Q'. (`The number of stars` relates in this way to *How many stars are there?*) That the area of overlap is non-empty means that S is true in α about $\mathtt{m}'$— in the weak sense that is not closed under conjunction. If the area of overlap contains moreover all the best worlds in α's $\mathtt{m}'$-cell, then S is true in α about $\mathtt{m}'$ in the strong sense that is closed under conjunction. The trick is to find a question whereof these are interesting things to know. Pérez-Carballo [2013] and Pérez Carballo [2014] contain illuminating discussions of these matters. See also Yalcin [2011]

Roman Frigg and Stephan Hartmann. Scientific models. In Sahotra Sarkar et al, editor, *The Philosophy of Science: An Encyclopedia, Vol. 2*. Routledge, 2005.

Axel Gelfert. Manipulative success and the unreal. *International Studies in the Philosophy of Science*, 17(3):245–263, 2003.

Peter Godfrey-Smith. Theories and models in metaphysics. *The Harvard Review of Philosophy*, 14(1):4–19, 2006.

Peter Godfrey-Smith. Models and fictions in science. *Philosophical studies*, 143(1):101–116, 2009.

Richard IG Hughes. Models and representation. *Philosophy of science*, pages S325–S336, 1997.

Angelika Kratzer. What 'must'and 'can'must and can mean. *Linguistics and Philosophy*, 1(3): 337–355, 1977.

Mary Leng. *Mathematics and Reality*. OUP Oxford, 2010.

Richard Levins. The strategy of model building in population biology. *American scientist*, pages 421–431, 1966.

Arnon Levy. Models, fictions, and realism: Two packages. *Philosophy of Science*, 79(5):738–748, 2012.

Arnon Levy. Modeling without models. *Philosophical Studies*, 172(3):781–798, 2015.

David Lewis. Statements partly about observation. In *Papers in philosophical logic*. Cambridge University Press, 1988.

Margaret Morrison and Mary S Morgan. Models as mediating instruments. *Ideas in Context*, 52:10–37, 1999.

Adam Morton. Mathematical models: questions of trustworthiness. *British Journal for the Philosophy of Science*, pages 659–674, 1993.

David Papineau. Mathematical fictionalism. *International Studies in the Philosophy of Science*, 2:151–174, Spr 88 1988. English.

LA Paul. Metaphysics as modeling: the handmaiden's tale. *Philosophical Studies*, 160(1):1–29, 2012.

Alejandro Pérez-Carballo. Good questions. unpublished, 2013.

Alejandro Pérez Carballo. Structuring logical space. *Philosophy and Phenomenological Research*, 2014.

Hilary Putnam. Models and reality. *Journal of Symbolic Logic*, 45(3):464–482, 1980.

Thomas C Schelling. *Micromotives and macrobehavior*. WW Norton & Company, 2006.

Stewart Shapiro. Mathematics and reality. *Philosophy of Science*, pages 523–548, 1983.

Burton Silver and Heather Busch. *Why cats paint: A theory of feline aesthetics*. Random House LLC, 2006.

Adam Toon. The ontology of theoretical modelling: models as make-believe. *Synthese*, 172(2): 301–315, 2010.

Bas C. Van Fraassen. *The Scientific Image*. Oxford University Press, USA, 1980.

Bas C Van Fraassen. Representation: The problem for structuralism. *Philosophy of Science*, 73 (5):536–547, 2006.

Kendall Walton. *Mimesis as make-believe: On the foundations of the representational arts*. Harvard University Press, 1990.

Kendall Walton. Metaphor and Prop Oriented Make-Believe. *European Journal of Philosophy*, 1(1):39–57, 1993.

Michael Weisberg. Three kinds of idealization. *The Journal of Philosophy*, pages 639–659, 2007.

Michael Weisberg. *Simulation and similarity: Using models to understand the world*. Oxford University Press, 2012.

Timothy Williamson. Model-building in philosophy. In Russell Blackford and Damien Broderick, editors, *Philosophy's Future: The Problem of Philosophical Progress*. Oxford: Wiley, 2016.

Stephen Yablo. Go figure: A path through fictionalism. *Midwest Studies in Philosophy*, 25(1): 72–102, 2001.

Stephen Yablo. Abstract objects: a case study. *Nous*, 36(s1):220–240, 2002.

Stephen Yablo. The myth of the seven. In Mark Kalderon, editor, *Fictionalism in Metaphysics*, pages 88–115. Oxford University Press, 2005.

Stephen Yablo. *Aboutness*. Princeton University Press, 2014.

Seth Yalcin. Figure and ground in logical space. unpublished, 2011.