

Scientific Theories*

Hans Halvorson

February 27, 2015

Abstract

Since the beginning of the 20th century, philosophers of science have asked, “what kind of thing is a scientific theory?” The logical positivists answered: a scientific theory is a mathematical theory, plus an empirical interpretation of that theory. Moreover, they assumed that a mathematical theory is specified by a set of axioms in a formal language. Later 20th century philosophers questioned this account, arguing instead that a scientific theory need not include a mathematical component; or that the mathematical component need not be specified by a set of axioms in a formal language. We survey various accounts of scientific theories entertained in the 20th century — removing some misconceptions, and clearing a path for future research. (Keywords: semantic view, syntactic view, received view, Carnap, van Fraassen, correspondence rules, category theory)

What is a scientific theory? Several philosophers have claimed that this question is *the* central philosophical question about science. Others claim still that the answer one gives to this question will fundamentally shape how one views science. In more recent years, however, some philosophers have become tired with this focus on theories — and they have suggested that we stop trying to answer this question. In this article, I will canvass and critically scrutinize the various answers that have been given to the question, “what is a scientific theory?” Then I will consider a recent argument against trying to answer this question. Finally, I will address the question of the utility of *formal* models of scientific theories.

*This is a preprint of an article forthcoming in *The Oxford Handbook of Philosophy of Science*. Please cite the published version.

1 The once-received view of theories

In the 1960s and 1970s, the vogue in philosophy of science was to identify problematic assumptions made by logical positivists — and to suggest that a better analysis of science would be possible once these assumptions were jettisoned. Of particular relevance for our discussion was the claim that the logical positivists viewed theories as “linguistic” or “syntactic” entities. Where did the 1960s philosophers get this idea? And is there any justice to their claim? [The references here are too numerous to list. Some of the most important include (Putnam 1962; Suppes 1964; Achinstein 1968; Suppe 1972; van Fraassen 1972; Suppe 1974).]

The story here is complicated by the history of formal logic. Recall that axiomatic systems of formal logic were common currency among philosophers of science in the 1920s, culminating in Rudolf Carnap’s *The Logical Syntax of Language* (Carnap 1934). At that time there was no such thing as formal semantics; instead, semantic investigations were considered to be a part of psychology or even of the dreaded metaphysics. Thus, when philosophers in the 1920s placed emphasis on “syntax,” they really meant to place emphasis on mathematical rigor. Indeed, what we now call “model theory” would almost certainly have been considered by Carnap et al. as a part of logical syntax. But more about this claim anon.

In any case, in his 1962 critique of the “received view” of scientific theories, Hilary Putnam describes the view as follows:

(RV) A scientific theory is a partially interpreted calculus.

What is meant by these notions? First of all, a “calculus” is a set of rules for manipulating symbols. What Putnam has in mind here is something like the “predicate calculus,” which involves: a set L of symbols (sometimes called a *signature*), a list of formation rules, and a list of transformation rules. The notion of “partial interpretation” is a bit more difficult to specify, a likely result of the fact that in the 1940s and 1950s, philosophers were still coming to terms with understanding model theory. In fact, in his critique of the received view, Putnam lists three ways of trying to understand partial interpretation, and rejects all three as inadequate.

The idea behind partial interpretation, however, is clear: some scientific theories rely heavily on various mathematical calculi, such as the theory of groups, or tensor calculus, or differential equations. But the statements of

mathematics don't, by themselves, say anything about the physical world. For example, Einstein's field equations

$$R_{ab} - \frac{1}{2}g_{ab} R = T_{ab},$$

will mean nothing to you unless you know that R_{ab} is supposed to represent spacetime curvature, etc.. Thus, Einstein's theory is more than just mathematical equations; it also includes certain claims about how those equations are linked to the world of our experience.

So, for the time being, it will suffice to think of "partial interpretation" as including at least an intended application of the formalism to the empirical world. We can then spell out RV further as follows:

(RV) A scientific theory consists of two things:

1. A formal system, including:
 - (a) Symbols;
 - (b) Formation rules; and
 - (c) Deduction rules.
2. Some use of this formal system to make claims about the physical world, and in particular, empirically ascertainable claims.

Perhaps the closest thing to an explicit assertion of RV is found in (Carnap 1939a, p. 193ff). [See also (Nagel 1961; Feigl 1970; Hempel 1970).] Before proceeding, note that Carnap himself was quite liberal about which kinds of formal systems would be permitted under the first heading. Unlike Quine, Carnap didn't have any problem with second-order quantification, intensional operators, non-classical logics, or infinitary logics. [When I need to be more precise, I will use $L_{\omega\omega}$ to indicate the first-order predicate calculus, where only finite conjunctions and disjunctions are permitted.]

The RV has been the focus of intense criticism from many different angles. In fact, it seems that between the years 1975 and 2010, beating up on RV was the favorite pastime of many philosophers of science. So what did they think was wrong with it?

The most obvious criticism of the RV, and one which Carnap himself anticipated, is that "scientific theories in the wild" rarely come as axiomatic systems. It is true that Carnap's proposal was based on some very special

cases, e.g. Einstein’s special theory of relativity, which admits of at least a partial axiomatization in $L_{\omega\omega}$. And it is not at all clear that other interesting scientific theories could be reconstructed in this way — not even Einstein’s general theory of relativity, nor quantum mechanics, not to speak of less formal theories such as evolutionary biology. (The problem with the former two theories is that they seem to require at least second-order quantification, e.g. in their use of topological structures.) So why did Carnap make this proposal when it so obviously doesn’t fit the data of scientific practice?

Here we must remember that Carnap had a peculiar idea about the objectives of philosophy. Starting with his book *The Logical Structure of the World* (Carnap 1928), Carnap aimed to provide a “rational reconstruction” of the knowledge produced by science. [For an illuminating discussion of this topic, see (Demopoulos 2007).] Carnap’s paradigm here, following in the footsteps of Russell, was the 19th century rigorization of mathematics afforded by symbolic logic and set theory. For example, just as 19th century mathematicians replaced the intuitive idea of a “continuous function” with a precise logically constructed counterpart, so Carnap wanted to replace the concepts of science with logically precise counterparts. In other words, Carnap saw the objective of philosophical investigation as providing a “nearest neighbor” of a scientific concept within the domain of rigorously defined concepts. For Carnap, if it was possible to replace individual scientific concepts with precise counterparts, then it was a worthy aim to formalize an entire domain of scientific knowledge.

Carnap’s ideas about “rational reconstruction” and of “explication” are worthy of a study in their own right. Suffice it to say for now that any serious discussion of Carnap’s views of scientific theories needs to consider the goals of rational reconstruction. [A nice discussion of these issues can be found in the introduction to (Suppe 1974).]

I’ve explained then why Carnap, at least, wanted to replace “theories in the wild” with formal counterparts. Many people now prefer to take a different approach altogether. However, in this article, I will mostly consider descendants of Carnap’s view — in particular, accounts of scientific theories that attempt to provide at least some formal precision to the notion.

But even among philosophers who agree with the idea of explicating “scientific theory,” there are still many objections to RV. For a rather comprehensive listing of purported difficulties with RV, see (Suppe 1974) and (Craver 2008). Rather than review all of these purported difficulties, I will focus on what I take to be misunderstandings.

1. *Fiction: RV treats scientific theories as linguistic entities.*

Fact: The RV gives the “theoretical definition” in terms of something that is often called a “formal language.” But a formal language is really not a language at all, since nobody reads or writes in a formal language. Indeed, one of the primary features of these so-called formal languages is that the symbols don’t have any meaning. Thus, we might as well stop talking about “formal language” and re-emphasize that we are talking about *structured sets*, namely, sets of symbols, sets of terms, sets of formulas, etc.. There is nothing intrinsically linguistic about this apparatus.

2. *Fiction: RV confuses theories with theory-formulations.*

Fact: To my knowledge, no advocate of RV ever claimed that the language-of-formulation was an essential characteristic of a theory. Rather, one and the same theory can be formulated in different languages. The failure to distinguish between theories and theory-formulations is simply a failure to understand the resources of symbolic logic. All that is needed to make this distinction is an appropriate notion of “equivalent formulations,” where two formulations are equivalent just in case they express the same theory. [For one reasonable account of equivalent theory formulations, see (Glymour 1971; Barrett and Halvorson 2015).]

The confusion here lies instead with the supposition that two distinct theory-formulations, in different languages, can correspond to the same class of models — a supposition that has been taken to support the semantic view of theories. This confusion will be unmasked in the subsequent section.

3. *Fiction: RV is inadequate because many interesting theories cannot be formulated in the first-order predicate calculus.*

Fact: I’ve already noted that Carnap, at least, was not committed to formulating theories in first-order logic. But even so, it’s not clear what is meant by saying that, “the theory T cannot be formulated in first-order logic,” when T hasn’t already been described as some sort of structured collection of mathematical objects.

Consider, for example, Einstein’s general theory of relativity (GTR). Is it possible to formulate GTR in a syntactic approach? First of all, this question has no definitive answer — at least not until GTR is

described in enough mathematical detail that its properties can be compared with properties of first-order axiomatizable theories. Moreover, it won't do to point out that, as it is standardly formulated, GTR involves second-order quantification (e.g. in its use of topological structure). For some theories formulated in second-order logic *also* admit a first-order axiomatization; or, in some cases, while the original theory might not admit a first-order axiomatization, there might be a similar replacement theory that does. One example here is the (mathematical) theory of topological spaces. While the definition of the class of topological spaces requires second-order quantification, the theory of "locales" can be axiomatized in (infinitary) first-order logic. And indeed, several mathematicians find the theory of locales to be a good replacement for the theory of topological spaces. [In fact, there is something like a first-order axiomatization of GTR, see (Reyes 2011).]

It's another question, of course, why a philosopher of science would want to try to axiomatize a theory when that would involve translating the theory into a completely alien framework. For example, I suspect that little insight would be gained by an axiomatization of evolutionary biology. But that's not surprising at all: evolutionary biology doesn't use abstract theoretical mathematics to the extent that theories of fundamental physics do.

Perhaps there is a simple solution to the supposed dilemma of whether philosophers ought to try to axiomatize theories: let actual science be our guide. Some sciences find axiomatization useful, and some do not. Accordingly, some philosophers of science should be concerned with axiomatizations, and some should not.

1.1 Correspondence rules

The most important criticism of RV regards the following related notions: correspondence rules, coordinative definitions, bridge laws, and partial interpretation. Each of these notions is meant to provide that additional element needed to differentiate empirical science (applied mathematics) from pure mathematics.

The notion of a coordinative definition emerged from late 19th century discussions of the application of geometry to the physical world. As was claimed by Henri Poincaré, the statements of pure mathematical geometry

have no intrinsic physical meaning — they are neither true nor false. For example, the claim that

(*P*) “The internal angles of a triangle sum to 180 degrees,”

says nothing about the physical world, at least until the speaker has an idea in mind about the physical referents of the words “line,” “triangle” etc.. A *coordinative definition*, then, is simply a way of picking out a class of physical things that correspond to the words or symbols of our mathematical formalism.

One example of a coordinative definition is a so-called operational definition. For example, Einstein defined two events to be *simultaneous* for an observer just in case that observer would visually register those events as occurring at the same time. Early logical positivists such as Hans Reichenbach took Einstein’s definition of simultaneity as a paradigm of good practice: taking a theoretical concept — such as simultaneity — and defining it in terms of simpler concepts. (Granted, concepts such as “seeing two images at the same time” are not as simple as they might at first seem!)

As the logical positivists came to rely more on symbolic logic for their explications, they attempted to explicate the notion of coordinative definitions within a logical framework. The key move here was to take the language L of a theory and to divide it into two parts: the observation language L_O , and the theoretical language L_P . By the late 1930s, the received view included this dichotomization of vocabulary, and efforts were focused on the question of how the terms in L_P “received meaning” or “received empirical content.”

It is the current author’s belief that Carnap and others erred in this simplistic method for specifying the empirical content of a scientific theory. However, I do *not* grant that the notion of empirical content cannot be specified syntactically, as has been suggested by van Fraassen (1980), among others. (For example, the distinction might be drawn among equivalence classes of formulas relative to interderivability in the theory T ; or the distinction might be drawn using many-sorted logic. The formal possibilities here seem hardly to have been explored.) Be that as it may, it was the simplistic way of specifying empirical content that was criticized by Putnam and others. With a devastating series of examples, Putnam showed that L_O terms can sometimes apply to unobservable objects, and L_P terms are sometimes used in observation reports. But let’s set those criticisms aside for the moment, and consider a second problem. Even if the vocabulary L could legitimately be divided into observational and theoretical components, there

remains the question of how the theoretical vocabulary ought to be related to the observation vocabulary.

In the years between 1925 and 1950, Carnap gradually loosened the restrictions he placed on the connection between theoretical (L_P) and observational (L_O) vocabulary. In the earliest years, Carnap wanted every theoretical term to be *explicitly defined* in terms of observation terms. That is, if $r(x)$ is a theoretical predicate, then there should be a sentence $\phi(x)$ in the observation language such that

$$T \vdash \forall x(r(x) \leftrightarrow \phi(x)).$$

That is, the theory T implies that a thing is r iff that thing is ϕ ; i.e. it provides a complete reduction of r to observational content. [Here Carnap was following Russell’s (1914) proposal to “construct” the physical world from sense data.]

However, by the mid 1930s, Carnap had become acutely aware that science freely uses theoretical terms that do not permit complete reduction to observation terms. The most notable case here is disposition terms, such as “ x is soluble.” The obvious definition,

$$x \text{ is soluble} \equiv \text{if } x \text{ is immersed, then } x \text{ dissolves} \quad ,$$

fails, because it entails that any object that is never immersed is soluble (see Carnap 1936, 1939b). In response to this issue, Carnap suggested that disposition terms must be connected to empirical terms by means of a certain sort of partial, or conditional, definition. From that point forward, efforts focused on two sorts of questions: were reduction sentences too conservative or too liberal? That is, are there legitimate scientific concepts that aren’t connected to empirical concepts by reduction sentences? Or, conversely, is the requirement of connectability via reduction sentence too permissive?

The final, most liberal, proposal about coordinative definitions seems to come from Hempel (1958). Here a theory T is simply required to include a set C of “correspondence rules” that tie the theoretical vocabulary to observational vocabulary. Around the same time, Carnap put forward the idea that theoretical terms are “partially interpreted” by means of their connection with observation statements. However, as pointed out by Putnam 1962, Carnap doesn’t provide any sort of precise account of this notion of partial interpretation. Indeed, Putnam argues that the notion doesn’t make any sense. Ironically, Putnam’s argument has been challenged by one of the strongest critics of the received view (Suppe 1971).

Thus, the so-called syntactic approach to theories was subjected to severe criticism, and was eventually abandoned. But I've given reason to think that a more sophisticated syntactic approach might be possible, and that such an approach would have all the advantages of the semantic approach to theories. In the next section, I'll also explain why I'm not convinced that semantic approaches have any intrinsic advantage over syntactic approaches. In fact, as I will explain in Section 5, the best versions of the syntactic and semantic approaches are formally dual to each other, and provide essentially the same picture of the structure of scientific theories.

2 The semantic view of theories

What I have been calling the “once received view” of theories is often called the “syntactic view” of theories — emphasizing that theories are formulated by means of logical syntax. According to Fred Suppe, the syntactic view of theories died in the late 1960s, after having met with an overwhelming number of objections in the previous two decades. At the time when Suppe wrote of the death of the syntactic view, it was unclear where philosophy of science would go. Several notable philosophers — such as Feyerabend and Hanson — wanted to push philosophy of science away from formal analyses of theories. However, others such as Patrick Suppes, Bas van Fraassen, and Fred Suppe saw formal resources for philosophy of science in other branches of mathematics, most particularly set theory and model theory. Roughly speaking, the “semantic view of theories” designates proposals to explicate theory-hood by means of the branch of mathematical logic called model theory. I will talk about the semantic view in this section, and I will discuss Suppes' set-theoretic account of theories in the following section.

Recall that the study of mathematical models (i.e. model theory) came alive in the mid 20th century with the work of Tarski and others. For philosophy of science, this advance was particularly significant, since the early positivists had banished semantical words such as “meaning,” or “reference” or “truth.” With the invention of formal semantics (i.e. model theory) these words were given precise explications.

However, philosophers of science were not all too quick to make use of model theory. One of the primary pioneers along these lines was the Dutch logician and philosopher Evert Beth. [A partial account of Beth's contribution to the development of the semantic view can be found in (van Fraassen

1970, 1972).] We will, however, proceed ahistorically and present the mature version of the semantic view.

Let L be a signature. An L -structure M (alternatively: an *interpretation* of L) is defined to be a set S , and an assignment of elements of L to relations or functions on Cartesian products of S . For example, if c is a constant symbol of L , then c^M is an element of S . As is described in any textbook of model theory, the interpretation M extends naturally to assign values to all L -terms, and then to all L -formulas, in particular, L -sentences. In fact, if an L -formula $\phi(x_1, \dots, x_n)$ has n free variables, then it will be assigned to a subset of the n -fold Cartesian product of S . As a special case, a sentence ϕ of L (which has zero free variables) is assigned to a subset of the singleton set, i.e. either the singleton set itself (in which case we say that “ ϕ is true in M ”), or the empty set (in which case we say that “ ϕ is false in M ”).

An L -structure is sometimes misleadingly called a “model.” This terminology is misleading, because the technically correct phrase is, “model of Σ ” where Σ is some set of sentences. In any case, we have the technical resources in place to state a preliminary version of the semantic view of theories:

(SV) A scientific theory is a class of L -structures for some language L .

Now, proponents of the semantic view will balk at SV for a couple of different reasons. First, semanticists stress that a scientific theory has two components:

1. A theoretical definition; and
2. A theoretical hypothesis.

The theoretical definition, roughly speaking, is intended to replace the first component of Carnap’s view of theories. That is, the theoretical definition is intended to specify some abstract mathematical object — the thing that will be used to do the representing. Then the theoretical hypothesis is some claim to the effect that some part of the world can be represented by the mathematical object given by the theoretical definition. So, to be clear, SV here is only intended to give one half of a theory, viz. the theoretical definition. I am not speaking yet about the theoretical hypothesis.

But proponents of the semantic view will balk for a second reason: SV makes reference to a language L . And one of the supposed benefits of the semantic view was to free us from the language dependence implied by the

syntactic view. So, how are we to modify SV in order to maintain the insight that a scientific theory is independent of the language in which it is formulated?

I will give two suggestions, the first of which I think cannot possibly succeed. The second suggestion works; but it shows that the semantic view actually has no advantage over the syntactic view in being “free from language dependence.”

How then to modify SV? The first suggestion is to formulate a notion of “mathematical structure” that makes no reference to a “language.” At first glance, it seems simple enough to do so. The paradigm case of a mathematical structure is supposed to be an ordered n -tuple $\langle S, R_1, \dots, R_n \rangle$, where S is a set, and R_1, \dots, R_n are relations on S . [This notion of mathematical structure follows Bourbaki (1970).] Consider, for example, the proposal made by Lisa Lloyd:

In our discussion, a *model* is not such an interpretation [i.e. not an L -structure], matching statements to a set of objects which bear certain relations among themselves, but the set of objects itself. That is, models should be understood as structures; in the cases we shall be discussing, they are mathematical structures, i.e., a set of mathematical objects standing in certain mathematically representable relations. (Lloyd 1984, p. 30)

But this proposal is incoherent. Let a be an arbitrary set, and consider the following purported example of a mathematical structure:

$$M = \langle \{a, b, \langle a, a \rangle\}, \{\langle a, a \rangle\} \rangle.$$

That is, the base set S consists of three elements $a, b, \langle a, a \rangle$, and the indicated structure is the singleton set containing $\langle a, a \rangle$. But what *is* that structure? Is that singleton set a monadic property? Or is that singleton a binary relation? (The former is a structure for a language L with a single unary predicate symbol; the latter is a structure for a language L' with a single binary relation symbol.) The simple fact is that in writing down M as an ordered n -tuple, we haven't really described a mathematical structure. Thus, a mathematical structure cannot simply be, “a set of mathematical objects standing in certain mathematically representable relations.”

To press the point further, consider another purported mathematical structure:

$$N = \langle \{a, b, \langle a, b \rangle\}, \{\langle a, b \rangle\} \rangle.$$

Are M and N isomorphic structures? Once again, the answer is underdetermined. If M and N are supposed to be structures for a language L with a single unary predicate symbol, then the answer is Yes. If M and N are supposed to be structures for a language L' with a single binary relation symbol, then the answer is No.

Thus, it's not at all clear how SV is supposed to provide a "language-free" account of theories. The key, I suggest, is to define a reasonable notion of equivalence of theory-formulations — a notion that allows one and the same theory to be formulated in different languages. But that same stratagem is available for a syntactic view of theories. Thus, "language independence" is not a genuine advantage of the semantic view of theories as against the syntactic view of theories. If the semantic view has some advantages, then they must lie elsewhere.

What then *are* the purported advantages of the semantic view of theories? What is supposed to recommend the semantic view? Here I will enumerate some of the advantages that have been claimed for this view, and provide some critical commentary.

1. *Claim: Scientists often work with heterogeneous collections of models that aren't all models of a single syntactically-formulated theory.*

This claim may well be true, but scientists engage in a lot of activities that don't involve constructing or evaluating theories. It seems that what might be suggested here is to consider the *ansatz* that the primary objective of science is the construction and use of models. I myself am loath to jump on the bandwagon with this assumption. e.g. mathematical physicists don't actually spend much time building models — they are busy proving theorems.

2. *Claim: Scientists often deal with collections of models that are not elementary classes, i.e. aren't the collection of models of some set of first-order sentences.*

This claim is strange, for it seems to indicate that scientists work with classes of L -structures (for some language L) that are not elementary classes (i.e. not the classes of models of a set of first-order sentences). I happen to know of no such example. Certainly, scientists work with classes of models that are not in any obvious sense elementary classes, but largely because they haven't been given a precise mathematical definition.

What about mathematical structures like Hilbert space, which are used in quantum mechanics? Isn't it obvious that the theory of Hilbert spaces is not elementary, i.e. there's no set Σ of first-order axioms such that the models of Σ are precisely the Hilbert spaces?

The problem mentioned above is still present. Although Hilbert spaces are fully legitimate citizens of the set-theoretic universe, the class of Hilbert spaces is not in any obvious way a class of L -structures for some language L .

What then are we to do in this case? Quantum mechanics is formulated in terms of Hilbert spaces, and so physics needs Hilbert spaces. If the received view can't countenance Hilbert spaces, then so much the worse for the received view, right?

I grant the legitimacy of this worry. However, there is a problem here not just for the received view, but for any philosopher looking at quantum mechanics (QM): while physicists use Hilbert spaces for QM, it is not at all clear what "theory" they are proposing when they do so. That is, it's not clear what assertions are made by QM. Or to restate the problem in a Quinean fashion, it's not clear what the domain of quantification in QM is.

One way to understand the task of "interpreting" quantum mechanics is finding the correct way to formulate the theory syntactically, i.e. finding the correct predicates, relations, and domain of quantification.

3. *Claim: Scientists often want to work with an intended model, but the syntactic approach always leaves open the possibility of unintended models.*

This point invokes the following well-known fact (the Löwenheim-Skølem theorem): for any theory T in $L_{\omega\omega}$ (assuming a countably infinite language), if T has a model of cardinality κ , then T has models of all smaller and larger infinite cardinalities. For example: if T is the first-order theory of the real numbers (say, in the language of ordered fields), then T has a countable model Q . But Q is not the model we intend to use if we believe that space has the structure of the continuum!

Once again, this "problem" can be simply dealt with by means of various technical stratagems, e.g. infinite conjunctions. But even if we remain in $L_{\omega\omega}$, it's not clear what this criticism was really meant to

show. On the one hand, the point could be that scientists need to discriminate between models that cannot be discriminated via first-order logic; e.g. a scientist might want to say that M is a good, or accurate model of some phenomenon, and N is *not* good or accurate, even though M and N are elementarily equivalent (i.e. they agree in the truth values they assign to first-order sentences). The key word here is “need.” I will grant that sometimes scientists’ preferences don’t respect elementary equivalence — that is, they might prefer one model over an elementarily equivalent model. But I’m not sure that this preference would have anything to do with them believing that their preferred model provides a more accurate representation of reality. They might well think that the differences between these models are irrelevant!

Suppose, however, that we decide that it *does* make a difference to us — i.e. that we want to be able to say: “the world might be like M , but it’s not like N , even though M and N are elementarily equivalent.” If we want to do that, must we adopt the semantic view of theories? Not necessarily: we could simply adopt a stronger language (say an infinitary logic, or second order logic) that allows us to discriminate between these models. If we claim that M has some feature that N lacks, then that’s because we are implicitly relying on a more expressive language than $L_{\omega\omega}$. In many paradigm cases, the distinctions can be drawn by means of second-order quantification, or even just with infinitary connectives and first-order quantification. The honest thing to do here would be to display our ontological commitments clearly by means of our choice of language.

4. *Claim: The semantic view is more faithful to scientific practice.*

“[A semantic approach] provides a characterization of physical theory which is more faithful to current practice in foundational research in the sciences than the familiar picture of of a partially interpreted axiomatic theory.” (van Fraassen 1970)

This criticism is partially dealt with by the fact that the syntactic view wasn’t supposed to provide a completely accurate description of what’s going on in science — it was supposed to provide an *idealized* picture. In the earliest years of logical positivism, Carnap explicitly described the relation of philosophy of science to science on analogy to the relation of mathematical physics to the physical world. The

point of this comparison is that mathematical physics makes idealizing assumptions so as to provide a tractable description that can be used to do some theoretical work (e.g. to prove theorems). In the same way, we can think of a syntactically formulated theory as an idealized version of a theory that provides *some*, but not complete, insight into the nature of that theory.

Perhaps the most helpful thing to say here is to echo a point made by Craver (2008) and also by Lutz (2015): the syntactic and semantic views are attempts to capture aspects of scientific theorizing, not the essence of scientific theorizing. Consequently, these two approaches need not be seen as competitors.

2.1 The problem of representation

Here's an initial puzzle about the semantic view of theories: if a theory is a collection of models, then what does it mean to *believe* a theory? After all, we know what it means to believe a collection of sentences; but what would it mean to believe a collection of models? At first glance, it seems that the semantic view of theories commits a basic category mistake.

Semanticists, however, are well aware of this issue, and they have a simple answer (albeit an answer that can be sophisticated in various ways).

(*B*) To believe a theory [which is represented by a class Σ of models], is to believe that the world is isomorphic to one of the models in Σ .

[See (van Fraassen 1980, pp. 68–69), (van Fraassen 2008, p. 309). Note that van Fraassen himself claims that “belief” is not typically the appropriate attitude to take toward a theory. The more appropriate attitude, he claims, is “acceptance,” which includes only belief in that theory’s empirical adequacy.] This idea seems simple enough . . . until you start to ask difficult questions about it. Indeed, there are a number of questions that might be raised about *B*.

First, “isomorphism” is a technical notion of a certain sort of mapping between mathematical structures. The physical world, however, is presumably not itself a mathematical structure. So what could be meant here by “isomorphism”? In response to this sort of question, some semanticists have suggested replacing the word “isomorphic” with the word “similar” (see Giere

1988, p. 83). Presumably the thought here is that using a less precise word will invite less scrutiny. In any case, a non-trivial amount of recent literature in philosophy of science has been devoted to trying to understand what this notion of similarity is supposed to be.

Second, the technical notion of isomorphism presupposes a specific *category* of mathematical structures. For example, consider the real numbers \mathbb{R} , and suppose that the intended domain of study X is supposed to be isomorphic to \mathbb{R} . What would that tell us about X ? Would it tell us something about the cardinality of X ? Would it also tell us something about the topological structure of X ? How about the order structure of X ? And what about the smooth (manifold) structure of X ? The point here is that \mathbb{R} belongs to several different categories — sets, groups, topological spaces, etc. — each of which has a different notion of isomorphism.

Third, and related to the second point, how are we to tell the difference between genuinely significant representational structure in a model and surplus, non-representational structure? For example, suppose that we represent the energy levels of a harmonic oscillator with the natural numbers \mathbb{N} . Well, which set did we mean by \mathbb{N} ? Did we mean the Zermelo ordinals, or the von Neumann ordinals, or yet some other set? These sets have different properties, e.g. the Zermelo ordinals are all singleton sets, the von Neumann ordinals are not. Does the “world” care whether we represent it with the Zermelo or von Neumann ordinals (cf Benacerraf 1965)?

The ironic thing, here, is that one the primary motivations for the semantic view of theories was to move away from problems that philosophers were creating for themselves, and back towards genuine problems generated within scientific practice. But the “problem of representation” — i.e. the question of how a mathematical model might be similar to the world — is precisely one of those problems generated internal to philosophical practice. That is, this problem was not generated internal to the practice of *empirical science*, but internal to the practice of *philosophical reflection on science*. To my eye, it is a sign of deterioration of a philosophical program when it raises more questions than it answers; and it appears that the semantic view has begun doing just that.

2.2 Criticisms of the semantic view

Critics of the syntactic, or received, view of theories have typically called for its rejection. In contrast, the semantic view of theories has mostly been sub-

jected to internal criticisms, with calls for certain modifications or reforms. Many of these discussions focus on the notion of “models.” While the earliest pioneers of semantic views were using the technical notion of “model” from mathematical logic, several philosophers of science have argued for a more inclusive notion of scientific models.

There have been a few external criticisms of the semantic view of theories; but most of these have been shown to rest on misunderstandings. For example, a classic criticism of the semantic view is that while one can believe in a theory, one cannot believe in a collection of models. But semanticists have been very clear that they see belief as involving the postulation of some notion of similarity or resemblance between one of the models and the intended domain of study.

One reason, however, for the paucity of criticism of the semantic view is that philosophers’ standards have changed — they no longer demand the same things from an account of theories that they demanded in the 1960s or 1970s. Recall, for example, that Putnam demanded that the received view give a precise account of partial interpretation. When he couldn’t find such an account, he concluded that scientific theories could not be what the received view said they were. The semantic view seems not to have been subjected to such high demands.

Consider, for example, the claim that the received view of theories entails that theories are language-dependent entities, whereas the semantic view of theories treats theories as language-independent. This comparison is based on a misunderstanding: classes of models *do* depend on a choice of language. Consider, for example, the following question: what is the class of groups? One might say that groups are ordered quadruples $\langle G, \circ, e, i \rangle$ where \circ is a binary function on G , i is a unary function on G , e is an element of G , etc.. Alternatively, one might say that groups are ordered triples $\langle G, \circ, i \rangle$, where \circ is a binary function on G , etc.. Or, for yet another distinct definition, one might say that a group is an ordered triple $\langle G, e, \circ \rangle$, where e is an element of G , etc.. Note that none of these classes are the same. For example, $\langle G, \circ, i \rangle$ cannot be a member of the last named class for the simple reason that i is a unary function on G . [For further discussion of this sort of criticism, see (Halvorson 2012, 2013; Glymour 2013; van Fraassen 2014).]

3 The set-theoretic view of theories

Beginning in the 1950s and 60s, Patrick Suppes developed a distinctive view of scientific theories as set-theoretic structures. In one sense, Suppes' view is a semantic view, insofar as mathematical semantics involves looking for structures inside the universe of sets. However, Suppes' approach differs in emphasis from the semantic view. Suppes doesn't talk about *models* but about *set-theoretic predicates*. He says,

To axiomatize a theory is to define a predicate in terms of the notions of set theory. (Suppes 1999, p. 249)

Recall that Zermelo-Frankel (ZF) set theory is a first-order theory in a language with a single binary relation symbol \in , where $x \in y$ intuitively means that the set x is an element of the set y . Following the typical custom, I'll make free use of the definable symbols \emptyset (empty set), \subseteq (subset inclusion), and ordered n -tuples such as $\langle x_1, \dots, x_n \rangle$.

In short, a set-theoretic predicate is just an open formula $\Phi(x)$ in the language of set theory. Thus, for example, the predicate “ x is an ordered pair” can be defined by means of the formula

$$\Phi(x) \equiv \exists y \exists z (x = \langle y, z \rangle).$$

Similarly, the predicate “ y is a function from u to v ” can be defined by means of a formula saying that y is a subset of ordered pairs of elements from u, v such that each element of u is paired with at most one element of v .

These set-theoretic predicates also allow us to define more complicated mathematical structures. For example, we can say that a set x is a “group” just in case $x = \langle y, z \rangle$, where y is a set, and z is a function from $y \times y$ to y satisfying certain properties. The result would be a rather complicated set-theoretic formula $\Gamma(x)$, which is satisfied by all and only those ordered pairs that are groups in the intuitive sense.

What are the advantages of the set-theoretic predicate (STP) approach? First of all, it is more powerful than the syntactic approach — at least if the latter is restricted to first-order logic. On the one hand, there are set-theoretic predicates for second-order structures, such as topological spaces. On the other hand, set-theoretic predicates can select intended models. For example, there is a set-theoretic predicate $\Phi(x)$ that holds of x just in case x is countably infinite.

Thus, it seems that the models of almost any scientific theory could be picked out by means of an STP. But just because you *can* do something, doesn't mean that you should do it. What advantage is there to formulating a scientific theory as an STP? Do such formulations provide some insight that we were looking for? Do they answer questions about theories? Moreover, what is the success record of Suppes' proposal?

Let's consider how Suppes' proposal fares in answering the sorts of questions philosophers of science might have about theories.

1. *Is the set-theoretic approach of any use for understanding the relations between models of a single theory?*

One advantage of a model-theoretic approach to theories is that model theory provides uniform definitions of the notions of *embeddings* and *isomorphisms* between models. e.g. if M, N are L -structures, then a map $j : M \rightarrow N$ is called an *embedding* just in case it "preserves" the interpretation of all non-logical vocabulary. Similarly, $j : M \rightarrow N$ is an *isomorphism* if it's a bijection that preserves the interpretation of the non-logical vocabulary. This single definition of isomorphism of L -structures then specializes to give the standard notion of isomorphism for most familiar mathematical structures such as groups, rings, vector spaces, etc.. It is this feature of model theory that makes it a fruitful mathematical discipline: it generalizes notions used in a several different branches of mathematics.

Now if A and B both satisfy a set-theoretic predicate $\Gamma(x)$, then when is A embeddable in B , and when are A and B isomorphic? A cursory scan of the literature shows (surprisingly!) that no such definitions have been proposed.

2. *Is the set-theoretic approach of any use for answering questions about relations between theories?*

Under what conditions should we say that two set-theoretical predicates $\Phi(x)$ and $\Psi(x)$ describe *equivalent* theories? For example, according to the folklore in mathematical physics, Hamiltonian and Lagrangian mechanics are equivalent theories — although this claim has recently been contested by philosophers (see North 2009). Suppose then that we formulated a set-theoretic predicate $H(x)$ for Hamiltonian mechanics, and another $L(x)$ for Lagrangian mechanics. Could these set-theoretic

formulations help us clarify the question of whether these two theories are equivalent?

Similarly, under what conditions should we say that the theory described by $\Phi(x)$ is *reducible to* the theory described by $\Psi(x)$? For example, is thermodynamics reducible to statistical mechanics?

For a rather sophisticated attempt to answer these questions, see (Pearce 1985), which develops ideas from Sneed and Stegmüller. Length considerations will not permit me to engage directly with these proposals. What's more, I have no intention of denigrating this difficult work. However, a cursory glance at this account indicates that it requires translating theories into a language that will be foreign to most working scientists, even those in the exact sciences. The problem here is essentially information overload: giving a set-theoretic predicate for a theory often means giving more information than a working scientist needs.

It is sometimes touted as a virtue of the set-theoretic approach that predicates in the language of set theory can be used to pick out intended models (and to rule out those unintended models, such as end extensions of Peano arithmetic). But this advantage comes at a price: there are also predicates of the language of set-theory to which scientists would show complete indifference. Consider an example: there are many sets that satisfy the set-theoretic predicate, “ x has the structure of the natural numbers.” Moreover, these different sets have different set-theoretic properties; e.g. one might contain the empty set, and another might not. That is, there is a set-theoretic predicate $\Phi(x)$ such that $\Phi(M)$ but $\neg\Phi(N)$, where M and N are sets that both instantiate the structure of the natural numbers. Thus, while set-theoretic predicates can be used to rule out unintended models, they also seem too fine-grained for the purposes of empirical science.

The criticisms I've made here of Suppes' approach can be considered as complementary to those made by Truesdell (1984). In both cases, the worry is that the set-theoretic approach requires translating scientific theories out of their natural idiom, and into the philosopher's preferred foundational language, viz. set theory. Obviously, set theory plays a privileged role in the foundations of mathematics. For example, we call a mathematical theory “consistent” just in case it has a set-theoretic model, in other words, just in case it can be translated into set theory. But the goal of philosophy of science

today is not typically to discover whether scientific theories are consistent in this sense; the goal, more often, is to see how scientific theories work — to see the inferential relations that they encode, to see the strategies for model building, etc.. — and how they are related to each other. The set-theoretic approach seems to provide little insight into those features of science that are most interesting to philosophers.

4 Flat versus structured views of theories

Typically the syntactic vs. semantic debate is taken to be the central question in the discussion of the structure of scientific theories. I maintain, however, that it's a distraction from a more pressing question. The more pressing question is whether scientific theories are “flat” or whether they have “structure.” Let me explain what I mean by this.

The syntactic view of theories is usually formulated as follows:

A theory is a *set* of sentences.

This formulation provides a *flat* view: a theory consists of a collection of things, and not in any relations between those things, or structure on those things. In contrast, a *structured* view of theories might say that a theory consists of both sentences and, e.g., inferential relations between those sentences.

The flat vs. structured distinction applies not just to the syntactic view of theories, but also to the semantic view of theories. A flat version of the semantic view might be formulated as:

A theory is a *set* (or *class*) of models.

In contrast, a structured version of the semantic view might say that a theory is a set of models *and* certain mappings between these models, such as elementary embeddings.

Both the syntactic and the semantic views of theories are typically presented as flat views. In the latter case, I suspect that the flat point of view is accidental. That is, most proponents of the semantic view are not ideologically committed to the claim that a theory is a bare set (or class) of models. I think they just haven't realized that there is an alternative account.

But in the case of the syntactic view, some philosophers have ideological commitments to a flat view — a commitment that derives from their rejection

of “intensional” concepts. The most notable case here is Quine. Quine’s criticism of the analytic-synthetic distinction can also be seen as a criticism of a structured view of theories. On a structured syntactic view of theories, what matters in a theory is not just the sentences it contains, but also the relations between the sentences (e.g. which sentences are logical consequences of which others). But in this case, commitment to a theory would involve claims about inferential relations, in particular, claims about which sentences are logical consequences of the empty set. In other words, a structured syntactic view of theories needs an analytic-synthetic distinction!

Quine’s powerful criticisms of the analytic-synthetic distinction raise worries for a structured syntactic picture of theories. But is all well with the unstructured, or flat, syntactic view? I maintain that the unstructured view has *severe* problems that have never been addressed. First of all, if theories are sets of sentences, then what is the criterion of equivalence between theories? A mathematically minded person will be tempted to say that between two *sets*, there is only one relevant condition of equivalence, namely equinumerosity. But certainly we don’t want to say that two theories are equivalent if they have the same number of sentences! Rather, if two theories are equivalent, then they should have some further structure in common. What structure should they have in common? I would suggest that, at the very least, equivalent theories ought to share the same inferential relations. But if that’s the case, then the content of a theory includes its inferential relations.

Similarly, Halvorson (2012) criticizes the flat semantic view of theories on the grounds that it would trivialize the notion of theoretical equivalence.

5 A category-theoretic approach to theories

As mentioned in Section 3, Suppes’ set-theoretic approach to theories was criticized by the historian of physics, Clifford Truesdell. Interestingly, a student of Truesdell’s became similarly agitated about the use of set-theoretic models for physical systems, seeing these models as obscuring the salient features of physical systems. This student was so firmly convinced that a new approach was needed that he devoted his career to developing an alternative foundation of mathematics, a “category-theoretic” foundation of mathematics.

Truesdell’s student was William Lawvere, who has gone on to be one of

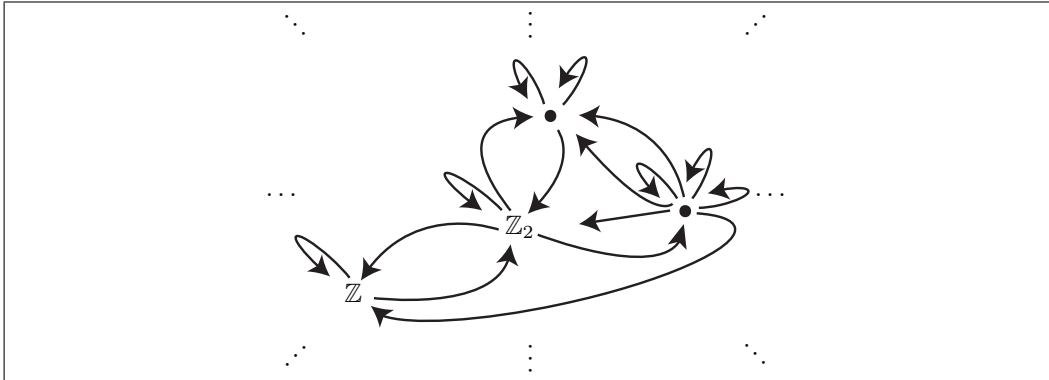


Figure 1: An impressionistic picture of one particular category. The nodes are the elements (or objects) of the category, and the lines are the paths (or arrows) between elements. We have labeled a couple of the nodes as familiar groups, since the collection of groups (objects) and group homomorphisms (arrows) is a paradigmatic category.

the most important developers of category theory (see Lawvere 2007). In retrospect, it is shocking that almost no philosophers of science followed the developments of Lawvere and his collaborators. To the best of my knowledge, there are at most a half dozen papers in philosophy of science that make use of category-theoretic tools.

So what is a category? One way of seeing it is that a category is a set where the elements can stand in a variety of relations to each other. In a bare set, two elements are either equal or unequal — there is no third way. But for two elements a, b in a category, there might be several different “paths” from a to b , some of which might be interpreted as showing that a and b are the same, and others showing that some other relation holds between them (see Figure 1).

Slightly more rigorously, a category C consists of a collection C_0 of “objects,” and a collection C_1 of “arrows” between these objects. The collection of arrows is equipped with a partial composition operation: if the head of one arrow meets the tail of another, then the two can be composed to give another arrow. Furthermore, for each object, there is an identity arrow going from that object to itself.

If you’re familiar with any abstract mathematical structures, then you are familiar with some categories. For example, the category **Grp** of groups has groups as objects and group homomorphisms as arrows. Similarly, the

category **Man** of manifolds has differential manifolds as objects and smooth mappings as arrows.

Some mathematicians and philosophers have made bold claims about category theory, e.g. that it should replace set theory as the foundation of mathematics. For the purposes of this essay, I won't need to take any position on that debate. Nonetheless, there are good reasons to think that category theory provides a better language than set theory for philosophers of science. Whereas set theory is best at plumbing the depths of individual mathematical structures (e.g. the real number line), category theory excels at considering large collections of mathematical structures and how they relate to each other. But isn't this point of view preferable for the philosopher of science? Isn't the philosopher of science's goal to see how the various components of a theory hang together, and to understand relations between different theories?

It was completely natural, and excusable, that when Suppes and collaborators thought of mathematical structures they thought of sets. After all, in the early 1950s, the best account of mathematical structure was that given by Bourbaki. But mathematics has evolved significantly over the past sixty years. According to current consensus in the mathematical community, the best account of mathematical structure is provided by category theory.

“The concept of mathematical structures was introduced by N. Bourbaki . . . The description which Bourbaki used was, unfortunately, rather clumsy. Simultaneously, a more abstract (and more convenient) theory of categories was introduced by S. Eilenberg and S. Mac Lane.” (Adámek 1983, ix)

Thus, if philosophers of science want to stay in touch with scientific practice, they need to think about the category-theoretic account of mathematical structure.

Categories have the capacity to represent theories as *structured* things — either structured syntactic things, or structured semantic things. From the syntactic point of view, a theory can be thought of as a category whose objects are sentences, and whose arrows are inferential relations between those sentences. From the semantic point of view, a theory can be thought of as a category whose objects are models, and whose arrows are mappings between those models. In both cases, category theory provides a particularly natural way to develop a structured view of theories.

Taking a category-theoretic view of theories has many advantages, and suggests many future projects for technically oriented philosophers of science.

1. Recall that flat views of theories have trouble providing a reasonable notion of theoretical equivalence. A category theoretic approach to theories fixes this problem. First of all, with regard to the syntactic point of view, the best account of theoretical equivalence is the notion of “having a common definitional extension,” as explained by Glymour (1971). It’s not difficult to see, however, that two theories have a common definitional extension just in case there are appropriate mappings between them that preserve inferential relations (see Barrett and Halvorson 2015).
2. To my knowledge, the semantic view of theories offers not resources for answering interesting questions about relations between different scientific theories. For example, are Hamiltonian and Lagrangian mechanics equivalent theories?

The reason that the semantic view cannot answer this questions is that if theories are treated as flat, structureless classes of models, then there is nothing interesting to say about relations between theories. It’s only when theories are treated as structured things that there are interesting mathematical questions about equivalence, reduction, and other relations between theories.

Once again, category theory can help here. If theories are categories, then a mapping between theories is a mapping between categories, something known as a “functor.” And there are interesting properties of functors that look very much like the relations between theories that interest philosophers of science. For example, a functor that is full, faithful, and essentially surjective looks like a good candidate for an equivalence of theories.

3. Some philosophers have claimed that there is no significant difference between the syntactic and semantic views of theories (see Friedman 1982; Halvorson 2012); or, at least, that in typical cases we can freely switch back and forth between syntactic and semantic representations of theories. In fact, category theory provides the means to make this claim rigorous.

First of all, given a syntactically formulated theory T , we can construct a class $M(T)$ of models. In fact, $M(T)$ is more than just a class — it is a category, since there is a natural notion of arrows between models,

namely elementary embeddings. Thus, M can in fact be thought of as a mapping from syntactically formulated theories to semantically formulated theories.

Is there a map N going in the reverse direction? Clearly, it cannot be the case that for any class Σ of models, there is a corresponding syntactically formulated theory $N(\Sigma)$. What's more, it's clear that the input for N shouldn't be bare classes of models, for one and the same class of models can correspond to different syntactically formulated theories (see Halvorson 2012). Thus, the input for N should be *categories* of models.

In short, there are interesting technical questions about moving from categories of models to syntactically formulated theories. The idea here is that when a category is sufficiently “nice,” then it is the category of models of a syntactically formulated theory. [These issues are discussed at length in (Awodey and Forssell 2013) and (Makkai and Reyes 1977; Makkai 1987), among other places in the literature on categorical logic.]

6 The no-theory view

Some philosophers of science will think that we've been wasting a lot of time splitting hairs about whether theories are syntactic things, or semantic things, or something else. And obviously, I somewhat agree with them. As indicated in the previous section, the best versions of the syntactic and semantic views of theories are dual to each other, and both analyses are helpful in certain philosophical discussions. However, some philosophers would go further and say that we needn't ever use the word “theory” in our philosophical discussions of science. In this section, I intend to rebut this intriguing suggestion.

Rather than try to deal with every critique of “theory,” I'll zoom in on a recent discussion by Vickers (2013), in the context of trying to understand how scientists deal with inconsistencies. According to Vickers, whenever we might ask something about a certain theory, we can just ask the same question about a set of propositions: “Why not simply talk about sets of propositions?” (p. 28) At this stage, it should be clear what problems this suggestion might lead to. First of all, we have no direct access to propositions — we only have access to sentences that express those propositions. Thus, we

only know how to talk about sets of propositions by using sets of sentences. Second, why should we think that when a scientist puts forward a theory, she's only putting forward a *set* of propositions? Why not think that she means to articulate a *structured set* of propositions, i.e. a set of propositions with inferential relations between them? Thus, we see that Vickers' stance is not neutral on the major debates about scientific theories. It's not neutral about the semantic vs. syntactic point of view; and it's not neutral on flat vs. structured points of view.

This is not to say that we can't go a long way in philosophical reflection on science without answering the question, "what is a theory?" Indeed, Vickers' book is a paradigm of philosophical engagement with issues that are, or ought to be, of real concern to working scientists and to scientifically engaged laypersons. But it seems to me that Vickers doesn't practice what he preaches. He disclaims commitment to any theory of theories, while his actual arguments assume that scientists are rational creatures who enunciate propositions, and believe that there are inferential relations between these propositions. Thus, I believe that Vickers' work displays the fruitfulness of a structured syntactic view of theories.

7 Why a formal account of theories?

Clearly this article has leaned heavily towards formal accounts of theories. Other than the author's personal interests, are there good reasons for this sort of approach? Once again, we are faced here with the question of what philosophers of science are hoping to accomplish; and just as there are many different sciences with many different approaches to the world, I suggest that there are many different legitimate ways of doing philosophy of science. Some philosophers of science will be more interested in how theories develop over time, and (*pace* the German structuralist school), formal analyses have so far offered little insight on this topic. Other philosophers of science are interested in how theory is related to experiment; and the more accurate and detailed an account of real-world phenomena we give, the more difficult it will be to illuminate issues via an abstract mathematical representation.

Thus, if a formal approach to scientific theories has utility, it has a limited sort of utility — in precisely the same way that mathematical physics has a limited sort of utility. Mathematical physics represents a sort of limiting case of scientific inquiry, where it is hoped that pure mathematical reasoning

can provide insight into the workings of nature. In the same way, formal approaches to scientific theories might be considered as limiting cases of philosophy of science, where it is hoped that pure mathematical reasoning can provide insight into our implicit ideals of scientific reasoning, the relations between theories, etc.. For some philosophers, this sort of enterprise will be seen as not only fun, but also illuminating. But just as there are scientists who don't like mathematics, it can be expected that there will be philosophers of science who don't like formal accounts of theories.

8 Further reading

For another perspective on scientific theories, see (Craver 2008), especially its account of mechanistic models.

For a historically nuanced look at the logical empiricists' account of theories, see (Mormann 2007). The most elaborate reconstruction of this approach is given in the introduction to (Suppe 1974). For a recent defense of some parts of the received view, see (Lutz 2012, 2014).

For a somewhat biased account of the development of the semantic view of theories, see (Suppe 2000). For another account of this approach, from one of its most important advocates, see (van Fraassen 1987). For discussion of the "problem of representation" in the semantic view of theories, see (Frigg 2006; Suárez 2010). For an argument that the syntactic-semantic question is a false dilemma, see (Lutz 2015).

For a detailed look at Suppes' set-theoretic approach, see (Suppes 2002). For a comparison of semantic and set-theoretic views of theories, see (Przelecki 1974; Lorenzano 2013).

Acknowledgments: Thanks to Thomas Barrett and Robbie Hirsch for sharpening my views on this subject, and for feedback on an earlier draft. Thanks to Jeff Koperski for bringing Truesdell's critique of Suppes to my attention.

References

- Achinstein, Peter. 1968. *Concepts of science*. The Johns Hopkins Press.
- Adámek, Jiří. 1983. *Theory of mathematical structures*. Springer.

- Awodey, Steve, and Henrik Forssell. 2013. "First-order logical duality." *Annals of Pure and Applied Logic* 164 (3): 319–348.
- Barrett, Thomas, and Hans Halvorson. 2015. "Quine and Glymour on theoretical equivalence." Unpublished manuscript. <http://philsci-archive.pitt.edu/11341/>.
- Benacerraf, Paul. 1965. "What numbers could not be." *The Philosophical Review* 74:47–73.
- Bourbaki, Nicholas. 1970. *Théorie des Ensembles*. Hermann.
- Carnap, Rudolf. 1928. *Der logische Aufbau der Welt*. Springer Verlag.
- . 1934. *Logische Syntax der Sprache*. Springer Verlag.
- . 1936. "Testability and meaning." *Philosophy of science* 3 (4): 419–471.
- . 1939a. "Foundations of logic and mathematics." In *International Encyclopedia of Unified Science*, edited by Otto Neurath, Rudolf Carnap, and Charles Morris, 2:139–212. Chicago.
- . 1939b. "Logical foundations of the unity of science." In *International Encyclopedia of Unified Science*, edited by Otto Neurath, Rudolf Carnap, and Charles Morris, 1:42–62. Chicago.
- Craver, Carl F. 2008. "Structures of scientific theories." In *The Blackwell guide to the philosophy of science*, edited by Peter Machamer and Michael Silbertstein, 55–79. John Wiley & Sons.
- Demopoulos, William. 2007. "Carnap on the rational reconstruction of scientific theories." In *The Cambridge Companion to Carnap*, edited by Michael Friedman and Richard Creath, 248–272. Cambridge University Press.
- Feigl, Herbert. 1970. "The 'orthodox' view of theories: Remarks in defense as well as critique." In *Analyses of Theories and Methods of Physics and Psychology*, edited by Michael Radner and Stephen Winokur, 4:3–16. Minnesota Studies in the Philosophy of Science. University of Minnesota Press.
- Friedman, Michael. 1982. "Review of *The Scientific Image*." *Journal of Philosophy* 79 (5): 274–283.

- Frigg, Roman. 2006. "Scientific representation and the semantic view of theories." *Theoria* 21 (1): 49–65.
- Giere, RN. 1988. *Explaining science: a cognitive approach*. University of Chicago Press.
- Glymour, Clark. 1971. "Theoretical realism and theoretical equivalence." In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, edited by Roger C. Buck and Robert S. Cohen, 275–288. Springer.
- . 2013. "Theoretical equivalence and the semantic view of theories." *Philosophy of Science* 80 (2): 286–297.
- Halvorson, Hans. 2012. "What scientific theories could not be." *Philosophy of Science* 79 (2): 183–206.
- . 2013. "The semantic view, if plausible, is syntactic." *Philosophy of Science* 80 (3): 475–478.
- Hempel, Carl Gustav. 1958. "The theoretician's dilemma: A study in the logic of theory construction." In *Concepts, theories, and the mind-body problem*, edited by Herbert Feigl, Michael Scriven, and Grover Maxwell, 2:37–98. Minnesota Studies in Philosophy of Science. Minneapolis: University of Minnesota Press.
- . 1970. "On the "standard conception" of scientific theories." In *Analyses of Theories and Methods of Physics and Psychology*, edited by Michael Radner and Stephen Winokur. University of Minnesota Press.
- Lawvere, F. William. 2007. "Interview with F. William Lawvere" (December). <http://www.cim.pt/files/publications/b23www.pdf>.
- Lloyd, Elisabeth. 1984. "A semantic approach to the structure of evolutionary theory." PhD diss., Princeton University.
- Lorenzano, Pablo. 2013. "The semantic conception and the structuralist view of theories: A critique of Suppe's criticisms." *Studies in History and Philosophy of Science Part A* 44 (4): 600–607.
- Lutz, Sebastian. 2012. "On a straw man in the philosophy of science: a defense of the received view." *HOPOS: The Journal of the International Society for the History of Philosophy of Science* 2 (1): 77–120.

- Lutz, Sebastian. 2014. “What’s right with a syntactic approach to theories and models?” *Erkenntnis* 79 (8): 1475–1492.
- . 2015. “What was the syntax-semantics debate in the philosophy of science about?” Unpublished manuscript. <http://philsci-archive.pitt.edu/11346/>.
- Makkai, Michael. 1987. “Stone duality for first order logic.” *Advances in Mathematics* 65 (2): 97–170.
- Makkai, Michael, and Gonzalo E Reyes. 1977. *First order categorical logic*. Springer.
- Mormann, Thomas. 2007. “The structure of scientific theories in logical empiricism.” In *The Cambridge Companion to Logical Empiricism*, edited by Alan Richardson and Thomas Uebel, 136–162. Cambridge University Press.
- Nagel, Ernest. 1961. *The structure of science*. Harcourt, Brace & World.
- North, Jill. 2009. “The “structure” of physics: A case study.” *The Journal of Philosophy* 106:57–88.
- Pearce, David. 1985. *Translation, Reduction and Equivalence*. Frankfurt: Peter Lang.
- Przelecki, Marian. 1974. “A set theoretic versus a model theoretic approach to the logical structure of physical theories.” *Studia Logica* 33 (1): 91–105.
- Putnam, Hilary. 1962. “What theories are not.” In *Logic, Methodology and Philosophy of Science: Proceedings of the 1960 International Congress*, edited by Ernest Nagel, Patrick Suppes, and Alfred Tarski, 240–251. Stanford University Press.
- Reyes, Gonzalo. 2011. “A derivation of Einstein’s vacuum field equations.” In *Models, Logics, and Higher-dimensional Categories*, edited by Bradd Hart, 245–261. American Mathematical Society.
- Russell, Bertrand. 1914. “The relation of sense-data to physics.” *Scientia* 16:1–27.
- Suárez, Mauricio. 2010. “Scientific representation.” *Philosophy Compass* 5 (1): 91–101.

- Suppe, Frederick. 1971. "On partial interpretation." *The Journal of Philosophy* 68:57–76.
- . 1972. "What's wrong with the received view on the structure of scientific theories?" *Philosophy of Science* 39:1–19.
- . 1974. *The structure of scientific theories*. Urbana, Illinois: University of Illinois Press.
- . 2000. "Understanding scientific theories: An assessment of developments, 1969–1998." *Philosophy of Science*:S102–S115.
- Suppes, Patrick. 1964. "What is a scientific theory?" In *Philosophy of Science Today*, edited by Sydney Morgenbesser, 55–67. Basic Books.
- . 1999. *Introduction to logic*. Courier Dover Publications.
- . 2002. *Representation and invariance of scientific structures*. Center for the Study of Language & Information.
- Truesdell, Clifford. 1984. "Suppesian stews." In *An idiot's fugitive essays on science*, 503–579. New York: Springer.
- van Fraassen, Bas. 1970. "On the extension of Beth's semantics of physical theories." *Philosophy of science* 37:325–339.
- . 1972. "A formal approach to the philosophy of science." In *Paradigms and Paradoxes*, edited by Robert Colodny, 303–366. University of Pittsburgh Press.
- . 1980. *The scientific image*. Oxford University Press.
- . 1987. "The semantic approach to scientific theories." In *The process of science*, edited by N.J. Nersessian, 105–124. Springer.
- . 2008. *Scientific Representation: Paradoxes of Perspective*. Oxford University Press.
- . 2014. "One or two gentle remarks about Hans Halvorson's critique of the semantic view." *Philosophy of Science* 81 (2): 276–283.
- Vickers, Peter. 2013. *Understanding inconsistent science*. Oxford University Press.