Truth in Explanation

by

Bernhard Nickel

Submitted to the Department of Linguistics and Philosophy
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2005

Author . . . .
Department of Linguistics and Philosophy
September 2, 2005

Certified by . . . .
Edward J. Hall
Professor of Philosophy
Thesis Supervisor

Certified by . . . .
Robert C. Stalnaker
Laurence S. Rockefeller Professor
Thesis Supervisor

Certified by . . . .
Stephen Yablo
Professor of Philosophy
Thesis Supervisor

Accepted by . . . .
Department Chair
Professor Stephen Yablo

Truth in Explanation

by

Bernhard Nickel

Submitted to the Department of Linguistics and Philosophy
on September 2, 2005, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

ABSTRACT

My thesis consists of three papers on truth and explanations in science. Broadly, the question I ask is semantic. Should the best account of certain bits of our scientific practice focus on the concept of truth? More specifically, should the crucial distinctions between good and bad aspects of that practice be drawn in terms of truth? My thesis consists of three case studies: ceteris paribus laws in the special sciences, appeals to idealizations in the application of theories, and the analysis of explanations quite generally, exemplified in the asymmetry of explanation. In each case, prominent philosophers have argued that a proper treatment does not focus on truth. In each case, I argue that truth should play a central role. And in each case, the issue turns, at least in part, on the connection between the scientific practice in question and explanations.

Thesis Supervisor: Edward J. Hall
Title: Professor of Philosophy

Thesis Supervisor: Robert C. Stalnaker
Title: Laurence S. Rockefeller Professor

Thesis Supervisor: Stephen Yablo
Title: Professor of Philosophy

Acknowledgments

This thesis is a status report on work in progress. It is no more, because my teachers and peers have given me one of the great gifts one can receive in graduate school. By example, they have helped me make myself as intellectually honest as I could be.

I have had the great fortune of being around some very smart people, indeed. One's fellow graduate students play a huge role in shaping one's education, and MIT has a long tradition of recruiting the best in the field. I learned a lot from many of them. I'd especially like to thank Selim Berker, Jason Decker, Tyler Doggett, Andy Egan, Chris Robichaud, Adina Roskies, Eric Swanson, and Seth Yalcin.

I'd also like to thank my teachers, who spent a lot of time on me. I want to especially thank my committee, Ned Hall, Stephen Yablo, and Robert Stalnaker, who pushed me at several crucial junctures to follow out the ideas in this thesis further than I was inclined to do on my own. They struck a wonderful balance between giving me enough guidance and enough space.

Judy Thompson deserves her own paragraph. In many ways, Judy is the one who trained me. The first time I learned how to reconstruct an argument was in the Proseminar (the second stage of Ayer's argument from illusion). I learned a lot about how to write a paper from the ten or so drafts I wrote for her class on events, and which she went over with me, in painstaking detail, in her office. More office visits have followed, and few things in life still fill me with more apprehension than letting Judy see my work. Her insight is truly remarkable, and she simply will not tolerate work from me that does not live up to my potential. My work and I am much the better for it. I owe her much.

Finally, I'd like to thank Susanna Siegel. She is my most important interlocutor about everything. Without her, it's not just this thesis that would not exist. Its author would not, either.

Biographical Sketch

Bernhard Nickel was born on September 26th, 1975 in Wiesbaden, Germany. After receiving his High School diploma ("Abitur"), he left Germany behind and moved to the U.S. He received his B.A. in 1999 with highest honors from Cornell University, majoring in philosophy. His work is focused on the philosophy of science and language. He also has interests in the philosophy of mind and epistemology.

Contents

Paper 1

□

CP-Laws, Processes, and Genericity

1 Introduction

The special sciences, such as biology, psychology, and economics, counsel us to accept *ceteris paribus* (cp) laws, generalizations that are held true despite the existence of counterexamples to the corresponding universal generalization. It is true that lions have four legs, even though some lions do not. There is controversy about the acceptability, and indeed the meaningfulness, of these claims. This controversy is the topic of this paper.

To structure the debate, we can distinguish two aspects of any law. A law is law-like and a generalization.[1] So for any proposed law, we can ask what sort of generalization is expressed, and whether that generalization satisfies the criteria for being law-like. In many cases, the first question receives an obvious answer, while the second is the hard one. But in the case of cp-laws, it is the first question that has received the most attention, and that is the question I will focus on, as well. So from hereon, I will talk about *cp-generalizations*, leaving open their status of law.[2] Specifically, I will focus on sentences of the form *As are Bs*, such as 'ravens are black' and 'lions have four legs'.

But I am not interested in all sentences of this form. To narrow my focus further, I will mean by a cp-generalization all and only the sentences of this form that also have the following two features.

---

[1]The contrast between law-like generalizations and the rest can perhaps be illustrated with this pair of examples (the examples are taken from Salmon (1989). Goodman's discussion of the new riddle of induction in Goodman (1983) is intended to draw attention to the same problem.):

(a) There are no gold spheres with a diameter greater than 100,000km.

(b) There are no uranium spheres with a diameter greater than 100,000km.

Both generalizations (a) and (b) are true, but allegedly only (b) is a law. So while it is not controversial that there is a distinction to be drawn, *how* to do it, and *whether* the distinction is important, is highly controversial.

[2]I doubt that my focus on the generalization expressed distorts the discussion I am addressing. The problems most writers on this topic discuss never have to do with what distinguishes laws from merely accidental generalizations, and instead only concerns the fact that laws of the special sciences tolerate exceptions. In order to not muddy the waters any more than absolutely necessary, I will continue to focus on cp-generalizations. An exception to the focus on the generalization part of cp-laws is Rosenberg (2001a,b). He seems to argue that cp-generalization, even well-confirmed ones, fail some of the other criteria of lawlikeness. I will not discuss his arguments here. For a response to these arguments, see Lange (1995, 2004).

[EXCEPTIONS] The generalization *As are Bs* would not be refuted by (at least some of) the things that would be counterexamples to the corresponding universal generalization *All As are Bs*.

[OPEN-ENDEDNESS] If we wanted to list the irrelevant potential exceptions to the generalization, we would have to use expressions like 'normal' or 'all things equal.'

OPEN-ENDEDNESS comes to this. Given EXCEPTIONS and that the people who articulate cp-generalizations take at least some of them to be true as stated, all of the instances of a cp-generalization cannot be equally relevant to its truth. When we say that lions have four legs, we simply do not count lions with an amputated leg as relevant. Let us call those instances of the generalization (i.e., those lions) that matter to its truth the *relevant* ones. The irrelevant ones are the ones that do not matter, and OPEN-ENDEDNESS says that we can only state which these are by using words like 'normal'.

We *need* to resort to phrases such as 'normal' because the irrelevant instances are sundry and miscellaneous: they form a thoroughly heterogeneous class. Hence, trying to spell out the generalization by adding a phrase beginning with 'unless' seems hopeless. For any list of acceptable exceptions we append in an 'unless'-clause, we can come up with further exceptions not yet included. Sooner or later, we will have to rely on phrases such as 'normal' to fill in the 'unless'-clause.

Consider (1).

(1) Lions have four legs.

Not all lions have four legs. Some have lost a leg in an accident. But it will not do to spell out (1) as (2).

(2) Lions have four legs unless they have lost some in an accident.

We have not covered all of the instances we would like to consider irrelevant, such as lions with birth defects. And it seems as if, for any way of spelling out (1) using an unless-clause that does not contain 'normal' or the like, we can come up with further irrelevant instances that are not yet covered by that clause. We can only say which instances of the generalization are irrelevant by using words like the ones mentioned in OPEN-ENDEDNESS.

All of the problems of cp-generalizations arise from OPEN-ENDEDNESS. Some theorists worry that OPEN-ENDEDNESS makes cp-generalizations trivial; some worry that it deprives them of truth-conditions; still others that it makes them unconfirmable. The triviality worry has so far received the most attention.

In this paper, I want to accomplish two things. First, I want to reorient the debate. I doubt that the triviality worry is a worry at all. This matters. As is the case quite generally, philosophical accounts of some phenomenon can be reductive or non-reductive. If the account is couched purely in terms that are better understood or "more fundamental", it is reductive. If not, it is non-reductive. The classic example of a reductive account is Goodman's response to his new riddle of induction. Here, appeals to entrenchment count as reductive, appeals to counterfactuals not.

Theorists who address themselves to the triviality worry accept that their account has to be reductive and that cp-generalizations are suspect until a reductive account is in hand. This is a problem, because it constrains our accounts illegitimately: we can make progress on more pressing concerns than triviality by giving a non-reductive account of cp-generalizations.

The reorientation of the debate away from triviality and away from reductive accounts is my first aim in this paper. The second aim is to give a non-reductive account that addresses an important puzzle about the confirmation of cp-generalizations. Here, too, my rejection of the triviality worry forces us to reconceive the problem. What is usually considered as the most pressing problem—that cp-generalizations cannot be confirmed or disconfirmed at all—does not get off the ground. Therefore, the puzzle I address differs from the one usually discussed. My response to the puzzle takes the form of spelling out more explicitly what cp-generalizations mean so that we can see more clearly how their confirmation might go. I will thus propose semantics for cp-generalizations.

In doing so, I am making an empirical claim. The semantics are correct for a bit of language that is used and meaningful independently of the stipulations of philosophers. I want to make good on this commitment by pointing out that cp-generalizations are a subset of a larger class of expressions we encounter in natural language, so-called *generics* and that my semantics are continuous with quite widely held views on generics that account for the data that linguists have drawn our attention to. This point sets my account apart from other proposals about the truth-conditions of cp-generalizations. Usually, authors just stipulate semantics and suggest that whatever problem concerning cp-generalizations they are interested in is solved. But the problem is only solved if they can show that their proposed semantics *really are* the correct semantics for cp-generalizations. Most do not.[3] Even when a theorist gives some motivation, as Lange (1995), the range of data accounted for is usually quite small. I hope to improve on this state by significantly increasing the range of data I consider.

Specifically, I will suggest that cp-generalizations depend for their truth on the normality of certain processes by which As can become or fail to become Bs. For example, 'ravens are black' is true just in case certain processes that bring ravens into existence are normal or characteristic for ravens. While discussions of generics in linguistics do not go so far as to posit processes in the interpretation of generics generally, or cp-generalizations in particular, my account is at least continuous with this discussion. What is more, the semantics I propose make it unmysterious how 'ravens are black' can depend for its truth only on facts about the normality of certain processes. If that is right, then there is no special problem

[3]See Fodor (1991), Glymour (2002), Lange (2000), and Pietroski and Rey (1995).

about confirming cp-generalizations. They are on a par with claims about the naturalness of natural kinds.

A crucial upshot of my response to the concern about confirmation is an objection to the empiricist hypothetico-deductive model of confirmation—what we may also call the instance-confirmation model of hypothesis testing. According to this model, a generalization is confirmed by observing an instance that conforms to it. That all ravens are black is confirmed by observing a black raven, for example. As both the paradox of the ravens and Goodman's 'grue' show, this model is much too simple.[4] Nonetheless we might think that there is something right about it. In order to get an adequate model of hypothesis testing, we just have to add further conditions to the instance-confirmation model.

I will argue that cp-generalizations form a class of generalizations to which the instance-confirmation model is simply inapplicable. Much of the interest in the model derives from the conviction that it applies to all or at least most of the generalizations scientists counsel us to accept. If it does not apply to cp-generalizations, that motivation is severely undermined. For all I say here, the model

[4]The paradox of the ravens is this. If we assume that the generalizations that all ravens are black receives the obvious first-order representation, a non-white non-raven is just as much an instance that conforms to the generalization as a black raven is. Thus, observing a white shoe, for example, should confirm that ravens are black. Clearly, it does not.

The problem raised by Goodman's 'grue' is that any set of observations is entailed by at least two mutually exclusive generalizations. But at most only one of these generalizations is intuitively confirmed by the data. Suppose 'grue' is defined thus: x is grue iff either (a) x is observed for the first time before noon tomorrow and x is green or (b) x is observed for the first time after noon tomorrow and x is blue. By the instance-confirmation model, a long string of observations of green emeralds up to the present should confirm both that all emeralds are green and that all emeralds are grue. Clearly, the observations only confirm the first.

may well apply to other generalizations, such as that all the coins in my pocket are made of copper. So I do not claim to refute the instance-confirmation model. I merely show that it is not nearly as important as its proponents would claim.[5]

The plan of the paper is as follows. I will begin by arguing that the triviality worry is not really worrisome (section 2). I then present one of the problems I take cp-generalizations to raise—the problem of how they are confirmed or disconfirmed by evidence (section 3). That is the problem I address in this paper. To do so, I present semantics for cp-generalizations (section 4). Finally, I apply the upshots of that discussion to the problem about confirmation (section 5) and tie up some loose ends (section 6).

2 Against Triviality

According to many authors, it is at least *prima facie* plausible that cp-generalizations are trivially true. They assume that to counter the worry, we must give a

[5]The problem about confirmation I try to identify is not the only problem one may raise about cp-generalizations. There is active debate about whether certain cp-generalizations about species that are pretheoretically acceptable are shown to be false by our best evolutionary theory. This is the debate about whether what is often known as "population-thinking" in evolutionary biology refutes a certain kind of essentialism. For arguments that it does, see Mayr (1994) and Sober (1994). For arguments that it does not, see Lange (1995).

Likewise, there is debate about whether cp-generalizations can serve as covering laws in a covering law theory of explanation. This debate occurs in the philosophy of history, for example. There, opponents of the covering law model of explanation argue that we at best only know cp-generalizations about history, and that these generalizations are not eligible to play the role of covering laws in covering law explanations. But since we do accept historical explanations, these theorists conclude that the covering law model is mistaken. On this point, see for example Dray (1957). I do not take a stand on these debates here, although I believe that my account has upshots for them, as well.

reductive account of their truth-conditions. Failing that, cp-generalizations are beyond the pale. Consider one prominent statement of the problem given by Pietroski and Rey (1995):

> There is a legitimate worry that appeals to cp-clauses render the nomic statements they modify somehow vacuous or unacceptably circular. In particular, if '$cp, F \Rightarrow G$' means merely that '$F \Rightarrow G$' is true in those circumstances in which there are no instances of F and not G, then 'cp-laws' look to be strictly tautologous.[6]

Pietroski and Rey do not think that the triviality worry carries the day. They defend semantics for cp-generalizations in the paper from which this quote is taken. But I want to object to the terms of the debate they take on board. First, there is a stark contrast between (1) and (3).

(1) Lions have four legs.

(3) Lions have fourteen legs.

We cannot bring ourselves to ever hear these two as even close in acceptability. The intuitive contrast rather suggests that (3) is false, not trivially true. And thus cp-generalizations cannot, in general, be trivially true.

[6]Pietroski and Rey (1995, p. 87). Woodward, who argues that there is a significant problem with cp-generalizations, cites this passage as stating the problem. See Woodward (2002, p. 308). Other authors who also take triviality to be the main problem include: Earman and Roberts (1999), Earman et al. (2002), Fodor (1991), Mott (1992), Schiffer (1991), Schurz (2001), and Schurz (2002).

One might think that the 'legitimate worry' mentioned by Pietroski and Rey has a different source. That cp-generalizations are trivially true is not so much a plausible hypothesis supported by our intuitive judgments of their acceptability, but rather forced on us by OPEN-ENDEDNESS. It is somehow plausible to infer triviality from OPEN-ENDEDNESS. But that just does not seem to be the case. We can certainly imagine situations in which (1) is false. For example, lions might evolve to have two legs. Two-legged lions in such a population would not be appropriately included in an unless-clause if we wanted to expand (1) along the lines of 'lions have four legs unless...'. This holds true of other generalizations, as well. It would be false that ravens are black if ravens had evolved to have the color scheme of parrots. That shows that cp-generalizations can be open-ended without being trivial.

One might argue for the restriction to reductive accounts more directly. CP-generalizations are part of a technical discourse whose technical terms—'ceteris paribus' in particular—need to be defined. That is to say, the technical terms require reductive accounts. There is a general objection that technical terms need not be explicitly definable in order to be used meaningfully in scientific inquiry. But more specific to the present topic, the phrase 'ceteris paribus' does not appear as a modifier in many cp-generalizations, and hence that expression should not be the target of analysis.[7] Rather, sentences such as 'ravens are black' and 'lions have

[7]I here agree with the Schiffer (1991, p. 10) and Woodward (2002, p. 305).

four legs' are the proper targets of analysis, and these sentences aren't couched in any distinctive vocabulary. And to suggest that there is an "implicit" *ceteris paribus* modifier is a substantive semantical hypothesis, not a starting point.

So the triviality worry is neither supported by intuitive acceptability judgments nor theoretical inference from OPEN-ENDEDNESS. We certainly should not take seriously the restriction to reductive accounts we would be bound by if we tried to respond to the triviality worry. And the failure of the direct argument suggests that we do not have independent reason to accept this stricture, either.

3 CP-Generalizations and Confirmation

While triviality is not a serious worry for cp-generalizations, their confirmation raises an important puzzle. But once we reject triviality as a worry, the usual problem about the confirmation of cp-generalizations disappears immediately. So let me begin with a representative statement of the standard problem, provided by Earman et al. (2002), *ERS* for short. They claim that cp-generalizations can neither be confirmed nor disconfirmed.

> In order for a hypothesis to be testable, it must lead us to some prediction. The prediction may be statistical in character, and in general it will depend on a set of auxiliary hypotheses. Even when these important qualifications have been added, CP law statements still fail to

> make any testable predictions. [...] our claim [that CP laws don't make testable predictions] is true only if none of the auxiliary hypotheses is the hypothesis that "other things are equal", or "there are no interferences". What if the auxiliaries *do* include the claim that other things are equal? Then either this auxiliary can be stated in a form that allows us to check whether it is true, or it can't. If it can, then the original CP law can be turned into a strict law by substituting the testable auxiliary for the CP clause. If it can't, then the prediction relies on an auxiliary hypothesis that cannot be tested itself. But it is generally, and rightly, presumed that auxiliary hypotheses must be testable if they are to be used in an honest test.[8]

The argument presented here fails, and the conclusion should be rejected in any case. First my objection to the argument. Let me just grant ERS that if we cannot be justified in believing that the instance we are observing is relevant, then we cannot be justified in believing the cp-generalization based on observing confirming instances. The point of contention concerns what it takes to be so justified. ERS suggest that we can justifiably believe that an instance is relevant only if we justifiably believe that there are no interferences—in a suitably broad sense of 'interferences.' The crucial requirement ERS impose is that we can justifiably believe *that* only if we can list all the possible interferences explicitly. OPEN-

[8] Earman et al. (2002, p. 293). Schurz gives the same argument for the impossibility of testing cp-generalizations experimentally in Schurz (2002, pp. 360-2).

ENDEDNESS says that this cannot be done.

But the explicit listing and checking of possible interferences is not generally required to have "an honest test." To see this, consider an uncontroversially testable hypothesis, such as that the Newtonian gravitational constant is such-and-such. One way to experimentally test this claim is to measure the mass of two objects, such as the Earth and a bowling ball, measure their distance from each other, and then release the bowling ball into a free fall. Finding out the ball's acceleration allows us to compute the value for the gravitational constant. Clearly, this is a good test only if gravity is the only force acting between the two bodies. In practice, this is a claim we need to make sure holds of our experimental set-up. The way we check is by considering all the forces our best scientific theories tell us about and make sure that these forces are absent or controlled for. But it does not follow logically from the fact that none of those forces act on the bowling ball, that none do at all.[9] That shows that we can be justified in believing that no interferences obtain even if we are in no position to list all possible interferences explicitly. So ERS' statement of the worry does not diagnose an intuitively compelling problem about the confirmation of cp-generalizations.

The problem also cannot be that cp-generalizations cannot be confirmed or disconfirmed, *simpliciter*. For many cp-generalizations, we can imagine situations

[9]Hempel has made the point that we need these kinds of provisos in order to test a theory in Hempel (1988). Indeed, opponents of cp-generalizations, such as ERS, grant that the fact that such provisos are needed does not raise any particular problems about testing theories: they even cite Hempel's argument, in order to show that there is no problem; see Earman et al. (2002, p. 285).

that would clearly falsify them.[10] If ravens had evolved to have the color scheme of parrots, it would be false that ravens are black. So if we had evidence that ravens really had so evolved, we would have disconfirmed 'ravens are black'. And if we have evidence that ravens have not so evolved, we have confirmed that claim.

If there is a problem about confirmation, it must arise from two observations. On the one hand, we can imagine situations that would falsify a cp-generalization we take to be true in the actual world. But on the other, there are no individual counter-instances of such a generalization. To see that this latter claim holds, think of a way, any way, of producing a non-black raven, such as a genetic mutation, an environmental toxin, a non-standard diet, or painting. The intuitive pull to say that any such raven is irrelevant to the generalization that ravens are black is very strong. That in turn suggests that no individual raven is a counter-instance to the generalization. But then it seems as if we obtain a situation that is not truly described by 'ravens are black' by collecting a large number of things that individually aren't counterinstances of the generalization. The puzzle is to say what the relationship is between individual instances of the generalization that fail to be counterinstances on their own and the situations we intuitively take to falsify the generalization.

Perhaps the simplest idea has it that cp-generalizations are claims about *most* of their instances: to say that lions have four legs is to say that most do. But this

[10] I want to exclude generalizations that are trivial by anybody's lights, such as 'ravens are ravens' or 'female ravens are ravens'.

will not do. Cutting a leg off every lion will not shake our conviction that lions have four legs. And to make the point even more forcefully, it is not even generally required that over the course of history, most instances need to conform to the cp-generalization. 'Mosquitos carry Plasmodia' is true, but it has never been true that most mosquitos carry Plasmodia.[11]

My proposal is instead that a cp-generalization *As are Bs* depends for its truth on facts about what makes an A a normal A. These change only when there are large-scale changes in the world, not when we consider merely a single instance of the generalization. This will also show that the hypothetico-deductive model of confirmation is inapplicable to cp-generalizations.

4 Semantics for Characterizing Sentences

4.1 Genericity in Natural Language

Genericity in natural language encompasses a range of different phenomena.[12] One kind, which I mention here only to set it aside, concerns reference to kinds (*genera*, hence the name 'generic'). Certain predicates, such as 'extinct,' 'rare,' and 'widespread' cannot be applied coherently to particular objects. Rather, they can only be applied to kinds as a whole.

[11]Plasmodia are the parasites that cause Malaria in humans.
[12]The range is, in fact, very large. For a lengthy list, see Krifka et al. (1995, p. 78).

(4) Bacteria are widespread.

(5) Dodos are extinct.

(6) Diamonds are rare.

Since the predicates that apply to kinds cannot apply to their instances, the predicates that apply to a kind do not apply to it *in virtue* of applying to some of its instances. In this respect, sentences that make reference to kinds differ from another kind of generic sentence, which I will call *characterizing sentences*.[13] These sentences express generalizations, and the predicate can also apply to the things generalized over.

(7) a. Lions have four legs.

b. Potatoes contain Vitamin C.

c. English verbs form their simple past tense by adding '-ed'.

All express generalizations. (7a) expresses a generalization over lions, (7b) a generalization over potatoes at times, and (7c) a generalization over English verbs.

The claim I'll be relying on is that cp-generalizations are a kind of characterizing sentence. The main justification comes from the observation that the sentences that paradigmatically express cp-generalizations pass all the tests linguists have

[13] I here follow, both in the terminology and the classification, Krifka et al. (1995, pp. 2-3).

proposed for detecting generic force. Aside from the fact that they tolerate exceptions, they share the following features with other characterizing sentences.[14]

- they are easily paraphrased by inserting 'normal' or 'normally';
- they usually cannot be paraphrased by using a progressive form; and
- they exhibit neither upward nor downward entailment.

In the introduction, I said that my semantics for cp-generalizations are continuous with a quite widely held view on generics. I now turn to this view.

4.2 The Semantics

As I said, I will assume that we are dealing with sentences of English of the form $\ulcorner As$ are $B\urcorner$, such as "ravens are black." Generic force can be conveyed with many other locutions, as well, such as singular nouns with either the definite or the indefinite article: 'the raven is black' and 'a raven is black' (at least on one of its readings).

A simplified version of the account I accept says that (8a) has the truth-conditions (8b).[15]

[14]For a longer discussion of these tests, see Krifka et al. (1995, pp. 8-14). I here simply assert that cp-generalizations pass these tests.

[15]For the idea of interpreting generics as having something to do with quantification and normality, see, for example, Asher and Morreau (1995), Dahl (1975), Farkas and Sugioka (1983), Heim (1982), Krifka et al. (1995), Lawler (1973). However, the particular proposals differ on the role of references to normality. For a defense of this particular proposal, as well as the required complications, see Nickel (forthcoming).

(8) a. As are Bs.

b. All normal As are Bs.

As it stands, this proposal is too simple. To name just two problems: right now, 'chickens lay eggs' entails 'chickens are hens,' and these semantics seem to commit me to countenancing ravens that are normal in all respects. Are these completely normal ravens male or female? But for present purposes, the simple version will do.

If we interpret (8a) as having the truth-conditions of (8b), we capture the fact that (8a) tolerates exceptions to the corresponding universal generalization because these instances are not normal. This proposal is intended to apply to all characterizing sentences, only some of which are cp-generalizations, because only some characterizing sentences have OPEN-ENDEDNESS. The difference between cp-generalizations and other characterizing sentences are captured by spelling out what normality amounts to in different ways. For now, it is just a place-holder.

5 Normality, Characteristic Processes, and the Concern about Confirmation

Simply appealing to normality does nothing to resolve the problem about confirmation I raised in the previous section. We still cannot see how we could arrive at a situation that falsifies a cp-generalization by collecting a number of things that individually aren't counter-instances. Now the problem is how it can be that we can arrive at a situation in which parrot-colored ravens are normal by collecting lots of these ravens that, by present lights, are not normal and hence not even what the generalization is about.

The key is to spell out normality as it appears in the interpretation of cp-generalizations. My proposal is that what makes something a normal instance of a given cp-generalization is that it goes through a certain kind of process. It will be easiest to present my proposal using an example, such as 'ravens are black.' All ravens come to be by going through some process or other. Among all of these processes, some are characteristic for ravens. A raven is normal just in case it has gone through one of the characteristic processes for ravens and their color. Thus, abnormal ravens are ones where no normal process got started, a normal process was interrupted, or its outcome was altered after the fact. More generally:

(9) x is a normal A in w iff the process that produced x is a characteristic

process for As in w.

Being characteristic is a world-bound property of processes: a process may be characteristic in one world but not another. This is brought out by thinking about ravens again. Let's suppose that in the actual world, a process issuing in a black raven is characteristic for ravens. But had ravens evolved to have the coloring of parrots, that very process would not have been characteristic for ravens. So we are left with this statement of the truth-conditions of cp-generalizations.

[TRUTH-CONDITIONS FOR CP-GENERALIZATIONS]

> ⌜As are Bs⌝ is true in w iff every x that is produced by a process characteristic for As in w is a B.

Let me now say a little more about processes. I identify processes with sets of events. They include the beginning of the process, its outcome, and significant events in between. Since the outcome of a given process is included among the identifying events, that a certain process is instantiated without interruption and without alteration after the fact guarantees that the outcome event obtains.

Here is the crucial fact about this interpretation of normality. Once we have settled which processes are characteristic for As in a world, we have settled which sentences of the form *As are Bs* are true in that world. Indeed, (10) holds necessarily.

(10) *A*s are *B*s iff all characteristic processes for *A*s issue in *B*s.

Suppose *A*s are *B*s. By the truth-conditions for cp-generalizations, that holds if and only if all normal *A*s are *B*s. That in turn holds if and only if all characteristic processes for *A*s issue in *B*s, by the interpretation of normality.[16]

5.1 Categorization By Process

On the view I'm defending here, a cp-generalization *As are Bs* reflects a division among *A*s into the normal ones and the rest. In order to defend this proposal, I need to show that this division plays a role in scientific practice, and that drawing the line by appeal to characteristic processes fits into our other theories of science. This will also enable me to say something about what makes a process characteristic.

It is a commonplace that we categorize in order to achieve inductive and explanatory success.[17] To say that the categories we deploy are explanatorily successful is to say at least that we can appeal to generalizations couched in terms of them in explanations. For example, our categorization allows us to confidently say that

[16]This derivation depends on the assumption that for every characteristic process there is an *A* that instantiates it. I discuss this assumption, and the modal commitments that come with it, in Nickel (forthcoming).

[17]The connection between categorization and induction has a long historical ancestry, as for example Mill (2002, Bk. I, Chp. vii). More recently, the connection has been made, inter alia, by Boyd (1991, 1999), Brandon (1997), Goodman (1983), Hacking (1991), Quine (1969), and Millikan (1999).

(11) It is easy to get sunburned while skiing because snow is white, and white snow reflects the sun's radiation.

The question for me thus becomes how categorization in terms of characteristic processes can help achieve this kind of explanatory success.

To make the connection, I want to appeal to a popular idea in the philosophy of science, what Boyd has called 'homeostatic property clusters,' or HPCs.[18] An HPC is a set of correlated properties where the correlation is ensured by some mechanism. The correlation may be maximal, but it need not. The crucial proposal I want to make is that we think of characteristic processes as the mechanisms that underwrite the correlation of properties in an HPC.

In some, and perhaps most cases, the characteristic process for *As are Bs* ensures that over the range of situations of interest to the scientific discipline, most As are Bs. Take the example (11) again. This is a good explanation only if most snow one encounters on a ski-slope is white. If the snow were black from pollution, (11) would be a terrible explanation. However, the characteristic process need not be such as to ensure that most As are Bs. For some explanatory aims, it is sufficient that the properties in an HPC overlap a little. Consider (12).

[18]See Boyd (1991, 1999), also Griffiths (1996), Hacking (1991), and Millikan (1999).

(12) Malaria is more common in tropical climates because mosquitos carry Plasmodia.

(12) is true, even though most mosquitos do not carry Plasmodia. Nonetheless, because there are so many mosquitos in tropical climates that even the minority of mosquitos that do carry them are a large number, malaria is more easily spread in climates hospitable to mosquitos.

So the non-negotiable connection between characteristic processes and HPCs is the use of cp-generalizations in explanations. Sometimes, this requires that if *As are Bs* is true, so is *Most As are Bs*, as in the example of snow. However, it need not. My suggestion is thus that characteristic processes are picked out by two factors conjointly, the explanatory aims of the science on the one hand, and on the other hand the features of the world that best serve these aims. Because of the second factor, the property of being characteristic is a world-bound property of processes. In different worlds, different processes serve the explanatory aims of a science.

On my view, then, we categorize twice over. We categorize once when we separate the As from all the other things, for example, when we distinguish ravens from non-ravens, lions from non-lions, and snow from other stuffs. We categorize again within each of these, separating the normal ravens from the other ravens, the normal lions form the other lions, and the normal snow from the rest. Why do we have to categorize twice over? I suggest that we categorize twice over because the two categorizations serve different ends. This is perhaps most dramatically illustrated in the case of species. The categorization into ravens and non-ravens is done for the purposes of evolutionary biology. As such, it is historical. But merely having a common historical origin, as all ravens do, is insufficient to ensure that all ravens also share causally relevant properties.[19] To capture a causally more homogeneous class, we need to focus on a subclass of ravens. My suggestion is that *that* is why we separate normal from non-normal ravens.

The appeal to HPCs that serve a science's explanatory success is more promising than a different tack one might take, suggested by the easy paraphrase of cp-generalizations as holding "so long as nothing interferes".[20] On this alternative, to find out what the characteristic processes for As are, take an A and see what happens to it in isolation. Then define these processes or suitably similar ones as the characteristic ones. But for most choices of A, nothing particularly interesting happens in isolation. All living creatures, for example, need sustenance. Without it, no creature would develop. The characteristic processes by which a living creature comes to be are ones in which it takes in some kinds of sustenance but not others. And in general, it does not seem as if As in isolation reveal very much.

So I've made two claims about characteristic processes: once we have settled which processes are characteristic for As, we have also settled which cp-genera-

[19]That historical categorization may well cross-cut causal categorization is pointed out by other theorists, as well. See Hacking (1991, 114-5) and Okasha (2002, 207-10).

[20]Something like this approach is suggested by Hall (2005).

lizations *As are Bs* are true. And because characteristic processes are determined by which HPCs are required to ensure the explanatory success of a science, which processes are characteristic for a given choice of A depends on the world. With these two claims in place, we can respond to the puzzle about confirmation.

5.2 Solving the Puzzle

In order to resolve the puzzle about confirmation, we need to say what the connection is between the instances of a cp-generalization that individually aren't counter-instances, such as parrot-colored ravens, and the situations that would falsify a generalization we consider true in the actual world, such as ravens evolving to have the color scheme of parrots.

On my semantics, our intuition that any individual raven that isn't black isn't relevant either tracks the semantic facts. There really aren't any relevant non-black ravens: once we have settled which processes are characteristic for ravens in a world, we have settled whether all of the ravens that go through the characteristic process are black. And given that in the actual world, all the characteristic processes for ravens issue in black ones, there are no ravens that are normal in this world but not black.

We can also see why large-scale changes in the world are enough to falsify a cp-generalization true in the actual world. Whether a particular process issuing in a raven of some kind is characteristic for ravens depends on large-scale patters in the world. So if we change these patterns, we change what is characteristic, and hence we may well change the truth-value of a cp-generalization. Put in slightly more intuitive terms, cp-generalizations are claims about categorization in the following sense. They are true just in case certain processes are categorized as characteristic, and hence just in case certain instances are categorized as normal. Small scale changes do not influence the categorization, large scale changes do.

If this response to the concern about confirmation is correct, cp-generalizations form a class of generalization to which the empiricist hypothetico-deductive model of hypothesis-confirmation does not apply. According to this model, a generalization is confirmed when one derives a prediction from it, perhaps with auxiliary hypotheses, and if that prediction is then verified. This model is either empty of all content, or clearly inapplicable to cp-generalizations.

Consider once again the generalization (13).

(13) Ravens are black.

Suppose that my semantics are correct. It seems clear that whether a certain process is characteristic for ravens is sensitive to facts about the world, just as whether gold is a natural kind is sensitive to some facts about the world. Call these facts in the case of ravens S, the supervenience base. Given suitable auxiliary hypotheses, we can deduce a statement of S from (13). These hypotheses need not be *ad hoc*, but they will be, at least in part, highly theoretically committed. We might then collect evidence about whether the facts in the supervenience base S really obtain. If the evidence says that they do, the generalization is confirmed. If the evidence says that they do not, it is disconfirmed. If we allow such an indirect evidential connection between the distribution of blackness in ravens and (13), then any kind of confirmation counts as an instance of the instance-confirmation model of hypothesis confirmation. Indeed, it is a completely trivial "model", since it says no more than that we confirm generalizations by finding evidence for them.

However, this does not seem to be what the empiricists had in mind when they proposed the model. To capture their spirit, we should place constraints on the auxiliary hypotheses we can use to derive the prediction to be tested. If these constraints rule out the auxiliary hypotheses needed to derive a statement of the supervenience base S from the generalization that ravens are black, then the instance-confirmation model has some content. But then it is clearly inapplicable to cp-generalizations. This does not show that the model is refuted, once it has been augmented with constraints on the auxiliary hypotheses. That is the wrong evaluative notion for this kind of model. Rather, these reflections show that the model is not nearly as widely applicable as empiricists might have hoped. It may still apply to other generalizations, such as that all the coins in my pocket are copper. But given that it does not apply to cp-generalizations, the model is much less central than its proponents would suggest.

My response to the concern about confirmation differs from the defense given by Fodor (1994). He accepts that in order to confirm a cp-generalization, we need to derive a prediction about a particular matter of fact from it, and that we can do so only with the help of the auxiliary hypothesis that "all things are equal." But Fodor parts ways with the opponents of cp-generalizations—such as Earman et al. (2002)—in saying that on particular occasions, we can be justified in believing that all things are, in fact, equal. As he puts it, that "all things are equal" is "epistemically accessible."[21] Fodor thus accepts that individual instances can refute or confirm a cp-generalization, and that the instance-confirmation model of hypothesis-testing applies to cp-generalizations. I deny both.

5.3 Elucidating Other Phenomena

I want to further support my proposal by pointing out some other features of cp-generalizations that are well accounted for if we posit such processes.

Open-Endedness By appealing to characteristic processes, we can say why cp-generalizations exhibit OPEN-ENDEDNESS. The phenomenon has two aspects. The list of irrelevant instances is open-ended, but some things definitely are not on the list. That is just the force of saying that cp-generalizations are not trivial. This latter aspect suggests that, in spite of the heterogeneity of the irrelevant instances of a given cp-generalization, there is some underlying principle that determines

[21] See Fodor (1994, p. 153).

whether an instance is relevant or not. By appealing to characteristic processes, we can account for both aspects of OPEN-ENDEDNESS.

For most cp-generalizations *As are Bs*, the characteristic processes issuing in *A*s are quite complex. The ways of interfering with that kind of process are legion. Indeed, it seems as if we are limited only by our imagination in the kinds of interferences we can come up with. That accounts for the heterogeneity of the irrelevant instances of a given cp-generalization.

The characteristic process also helps us to see what all of the irrelevant instances have in common: they are the result of interferences with these processes. This in turn sheds light on the feature of OPEN-ENDEDNESS itself. When I introduced it, I said that we have to use words like 'normal' to state which instances are irrelevant. At the time, I did not say what the force of *words like* 'normal' was. Now I can. Any way of spelling out the interferences with the characteristic processes will count as using words that are, in this connection, relevantly like 'normal.'

Finally, by appealing to characteristic processes, we can elucidate another aspect of open-endedness. For a given generalization *As are Bs*, even some *A*s that *are B*s are irrelevant to the generalization, such as an albino painted black. An albino is an instance in which the characteristic process never got off the ground, and hence does not count as normal, no matter what else happens to it.

My account of why cp-generalizations exhibit OPEN-ENDEDNESS is superior to the other explanations currently extant in the literature. These are suggested by Fodor (1994) and Earman and Roberts (1999). I will briefly discuss each.

Fodor says:

> Special science laws are unstrict not just de facto, but in principle. Specifically, they are characteristically "*heteronomic*": you can't convert them into strict laws by elaborating their antecedents. One reason why this is so is that special science laws typically fail in limiting conditions, or in the case of conditions where the idealizations presupposed by the science aren't approximated; and generally speaking, you have to go outside the vocabulary of the science to say what these conditions are. Old rivers meander, but not when somebody builds a levee. Notice that "levee" is not a *geological* term. (Neither, for that matter, is "somebody.")[22]

At least at first sight, this seems to be a non-sequitur. What is to be explained is why a cp-generalization cannot be turned into a strict law by stating 'unless'-clauses that do not contain something like 'normal'. Fodor's explanation only shows why cp-generalizations cannot be elaborated using a more restricted vocabulary, that of a particular special science. But that does nothing to explain why we cannot spell out what makes an instance relevant even without such strictures.

[22]Fodor (1994, p. 147n10). Emphasis in the original.

Moreover, it seems wrong to assume that cp-generalizations express idealizations. There are no frictionless planes, and so it is an idealization that on an idealized plane, the acceleration of a ball depends only on the angle of the plane and the force of gravity. But there are normal ravens who come by their color in the usual way.

Now, perhaps Fodor assumes that a claim like 'old rivers meander' means that all old rivers do, and it is therefore in some sense an idealization if we describe the actual world, in which some old rivers do not meander, as one in which they all do. But I have already argued that the generalization does not mean that. And what is more, even then it is not apt to call this an idealization. If I say "all bottles are on the table", meaning to convey that all bottles in the room are on the table, I am not idealizing, I am restricting my attention. So Fodor's analysis, if it is meant as an explanation of why cp-generalizations exhibit OPEN-ENDEDNESS, fails.[23]

Earman and Roberts (1999) also propose an explanation. At any rate, they announce that they seek to explain why generalizations in the special sciences must be cp-generalizations. Since on their view, something counts as a cp-generalization only if it exhibits OPEN-ENDEDNESS, we might expect an explanation of

[23] I state the conclusion conditionally, since the note I quoted at the beginning is followed by this remark: "I emphasize this point because it's sometimes supposed that heteronomicity is a proprietary feature of *intentional* laws qua intentional. Poppycock." (Fodor (1994, p. 147n10)) So perhaps Fodor only wants to establish the heteronomicity of cp-generalizations. However, in that case it becomes unclear why he starts the note by stating what I take as his explanandum in the text, that cp-generalizations exhibit OPEN-ENDEDNESS, or that they are "unstrict in principle", as he puts it.

this feature.

> if [the special sciences] are to propose any laws at all, then these will apparently have to be CETERIS PARIBUS laws.[24]

However, as their ensuing discussion makes clear, their argument establishes at most that for a given special science S and cp-generalization G of this science, one cannot state what is a relevant instance of G in the vocabulary of S.[25] That is to say, they establish the same conclusion as Fodor does in the passage cited above. And as I have argued above, that does nothing to explain why cp-generalizations exhibit OPEN-ENDEDNESS.

Commonalities within a Science I have suggested that the characteristicness of a process depends on the explanatory aims of a scientific discipline. The next two features of cp-generalizations are unsurprising if that's right.

Consider generalizations that are, intuitively, part of the same discipline, such as (14)-(16).

(14) Ravens are black.

(15) Tigers are orange with black stripes.

[24] Earman and Roberts (1999, p. 446).

[25] See Earman and Roberts (1999, pp. 447-8), citing Schiffer (1991) and Cartwright (1995).

(16) Dolphins are gray.

Once we know under what conditions a raven is irrelevant to (14), we also have a fairly good idea of the conditions under which a tiger is irrelevant to (15), and the conditions under which a dolphin is irrelevant to (16). Clearly, there is some systematicity to these generalizations. We can capture it by appeal to characteristic processes.

Specifically, the similarity concerns what sorts of influences count as derailing the characteristic process. For all three, human interaction is not part of the characteristic process. Likewise, painting is out. If I was right earlier in saying that characteristic processes are picked out by considering the explanatory aims of the science in question, we might expect not just that some generalizations share similar characteristic processes, but that all do. So appeal to characteristic processes might allow us to individuate scientific disciplines: two generalizations belong to the same discipline just in case the characteristic processes for both are similar.

Generalizations across the Disciplines When we pursue comparative zoology, human interference is not part of the characteristic processes. But when we engage in dog-breeding, it is. To that end, consider the generalization (17).

(17) Dobermans have floppy ears.

When we engage in comparative zoology, (17) is true, because human interference is not part of the characteristic process. But in the context of dog-breeding, it is part of the characteristic process for dobermans that their ears are cut to make them stand up. So depending on the area of inquiry, human interference may or may not be part of the characteristic process for dobermans.[26]

6 The Target Class of Sentences

I have so far only given a positive characterization of cp-generalizations. I now want to explicitly distinguish them both from other characterizing sentences, as well as from sentences that are sometimes lumped in with cp-generalizations in the literature. I do this to forestall objections that point to the fact that my account may not work for these other sentences. In each case, I'll argue that the sentences in question don't have the two features of EXCEPTION and OPEN-ENDEDNESS.

[EXCEPTIONS] The generalization would not be refuted by (at least some of) the things that would be counterexamples to the corresponding universal generalization.

[OPEN-ENDEDNESS] If we wanted to list the irrelevant potential exceptions to

[26] That cp-generalizations about species depend for their acceptability on the scientific context is a point made by Lange (1995), who suggests that, though they may be unacceptable in evolutionary biology, they are acceptable when we engage in zoology. Indeed, the more general point that different disciplines will deploy different schemes of categorization is quite common. See for example Hacking (1991) and Miller (2000).

the generalization, we would have to use expressions like ‘normal’ or ‘all things equal.’

6.1 Ascriptions of Dispositions

A form of sentence often counted as a cp-generalization, and hence as something that raises the same semantic problem, are ascriptions of dispositions. A common example is (18).[27]

(18) Metal bars expand when heated.

(18) ascribes a disposition to metal bars, the disposition to expand when heated. Which bars? All of them. Every metal bar has the disposition of expanding when heated. Sometimes, a bar manifests this disposition. Sometimes, it does not. But distinguishing between these different situations is not the same as distinguishing between relevant and irrelevant metal bars. Hence, they do not satisfy the condition EXCEPTION.

A slightly more loaded way of making this point is to consider mixed cases, such as (19).

(19) Dogs bark.

[27] See for example Lange (2000, chp. 6).

(19) ascribes a disposition, the disposition to bark under certain circumstances. But (19) does not ascribe this disposition to every dog. For example, the fact that dogs with a defective vocal tract fail to have the disposition does not falsify (19).

A way of paraphrasing (19) is as (20).

(20) Normal dogs normally bark.

Intuitively, 'normal' rules out dogs without vocal tracts; 'normally' ensures that (20) is compatible with the fact that sometimes, even normal dogs do not bark. The two occurrences of 'normal' are not redundant, so we should not expect the two modifiers to mean the same thing. It is possible that an account similar to mine can be made to work for ascriptions of dispositions. But this need not be the case, and my account is not committed to its being the case.

For some generalizations, it may not be immediately obvious whether they are cp-generalizations. One puzzling case concerns economics. Consider

(21) If the price of a good falls, demand for it rises.

It is tempting to think that this is a mixed case, akin to (19). On the one hand, there are kinds of goods to which the generalization is intuitively inapplicable, such as luxury goods, or goods that have completely inelastic demand.[28] On the

[28]Luxury goods might be disqualified, because part of the appeal of a true luxury good is its price. By owning one, one shows that one is the sort of person who can afford it; but if the price

other hand, even for goods that do not fall into any of these categories, there are situations in which even though their price drops, demand for it does not rise. That might happen, for example, if an alternative good drops in price even more, and so consumers shift their purchases to the alternative good. So (21) ascribes a disposition to such "normal" goods that is sometimes manifested, sometimes not.

However, I do not think that this distinction holds up under scrutiny. There is nothing intrinsic to, say, diamonds that makes them luxury goods. Rather, what makes diamonds a luxury good is that we are disposed to treat them as such. But nothing stands in the way of treating other things the same way, and ceasing to treat diamonds as luxury goods. At any rate, nothing in the nature of diamonds stands in this way. There are historical and social forces that do. Likewise, there is nothing intrinsic to water that makes it have inelastic demand. Rather, the inelasticity of the demand is due to how we are disposed to treat it. Granted, these dispositions are due to relatively immalleable aspects of human beings, but that is a far cry from saying that the inelasticity of the demand for water is due to some intrinsic features of water. That suggests that really, there is no distinction between "normal" goods like coffee and abnormal goods like water and diamonds.

Of course, my insistence on distinguishing between cp-generalizations and ascriptions of dispositions might strike some as missing the point. Ascriptions of

drops, the good can no longer play this role. So we might think that if the price of a luxury good drops, demand for it drops as well. A good is inelastic if a consumer's purchasing behavior is unaffected by the price of a good. Water is an inelastic good.

dispositions and cp-generalizations raise very similar problems. In both cases, a natural way to spell out the meaning of a given sentence is by making use of an unless-clause. And the list of things in the unless-clause is open-ended in very similar ways. In both cases, it is natural to cap the unless-clause by using 'unless something interferes.' So it would be nice if a unified treatment of cp-generalizations and ascriptions of dispositions were available. I share this view. But we should also expect that there are some differences between the two kinds of sentences, as is exhibited by the fact that 'dogs bark' is properly paraphrased with two occurrences of 'normal': 'normal dogs normally bark'. If a single account of ascriptions of dispositions and cp-generalizations were available, we should expect the two occurrences of 'normal' to be reducible to one. But they are not. So instead, we should expect accounts of the two kinds of sentences to be relevantly similar, without being quite the same.

6.2 IDEALIZATIONS

Consider a statement of a paradigmatic idealization, the ideal gas law.

[IDEAL GAS LAW] $PV = nRT$.[29]

The ideal gas law is true of every ideal gas, so it does not satisfy EXCEPTIONS. Moreover, suppose we appeal to the IDEAL GAS LAW in solving some physics

[29] Here, P is the pressure of the gas, V its volume, T the temperature, n the number of mol, R the universal gas constant.

problem, and we are told that the actual gas we are thinking about is not an ideal gas. The appropriate response is not to begin by distinguishing relevant from irrelevant bits of gas. So statements of idealizations do not exhibit OPEN-ENDEDNESS, either.[30] So statements of idealizations are not cp-generalizations. I thus agree with several authors that the problem of idealization is, at least semantically, importantly different from that posed by what I call cp-generalizations.[31] Other authors claim that idealizations fall under the heading of 'ceteris paribus' laws.[32] However, these authors usually do not give any clear criteria for what unifies the various claims they group under the heading of cp-laws. The mere fact that a claim would be false if read as a strict law about things in this world seems to suffice. They are then forced to distinguish those generalizations that express idealizations, such as the IDEAL GAS LAW, from others. So one way or another, everybody has to recognize that statements of idealizations deserve a different treatment from cp-generalizations.

[30]That is not to say that there is no interesting problem about how we use idealizations, or what makes a use good when it is. I address these questions in my paper "Truth and the Presence of Models."

[31]See Earman and Roberts (1999), Earman et al. (2002, p. 285), and Smith (2002).

[32]See Morreau (1999, p. 168), Pietroski and Rey (1995) and Schurz (2002, p. 352). Pietroski and Rey (1995) explicitly state that "the need for cp-laws stems from the need to idealize in a complex world."(p. 84).

6.3 THE LAWS OF PHYSICS

Some theorists have suggested that the laws of physics hold only *ceteris paribus*.[33] The reason is *not* that laws like Newton's have remained in use even after they have been superseded by better theories which treat Newton's laws as special cases. These theorists content that even had Newton's laws remained the last word, they would still only have held *ceteris paribus*. The easiest example is perhaps the law of gravity. It says that an object o_1 with some mass in the vicinity of another object o_2 experiences a force that is a function of the other object's mass and their distance. Now, if we take the force at issue to be the *total* force acting on o_1, this will often be false. A feather close to the Earth, for example, is also subject to considerable friction-forces. Hence, the law of gravity is a cp-law.

If it is true that laws of physics only hold *ceteris paribus*, my account would obviously be sunk. There is no characteristic process in the vicinity. But I, along with many theorists, want to resist the conclusion of the argument. The force at issue in the law of gravity is not the total force acting on the object, but merely the force *due to gravity*. So the laws of physics may still hold universally.[34]

[33]Most famously, Cartwright (1983) holds this view. See also Cartwright (1999a), as well as Giere (1988) and Giere (1999, chp. 5).

[34]For similar arguments, see Earman et al. (2002) and Smith (2002).

6.4 Probabilistic (Property) Causation

Some theorists have suggested that (22) is a cp-generalization.[35]

(22) Smoking causes cancer.

I take it that roughly, (22) says that smoking is a causal contributor to developing cancer. As a first pass, we can interpret that as saying that smoking raises the likelihood of developing cancer in every smoker. A smoker who does not develop cancer is not a counter-instance to this generalization, since it may be true that her cancer-risk was elevated because of her smoking. She just happened to be lucky. In that case, (22) is just a straightforward universal generalization, and a true one at that.

However, the first pass interpretation may be wrong. Some authors report the intuition that (22) would be true even if in some very small subgroup of smokers, smoking actually lowered the likelihood of cancer.[36] Suppose that this intuition is compelling. Then one might say that (22) is a cp-law after all, because only some subgroups of a given population are relevant. However, whether this is the case or not depends on how exactly we want to respond to the intuition I just mentioned. One response—the one favored by Dupré (1984)—is to say that (22) makes a purely statistical claim: it is true just in case in most of the population,

[35]See Earman et al. (2002, p. 295)

[36]Such as Dupré (1984).

smoking raises the likelihood of cancer. If that response is taken, (22) is not a cp-generalization, because it does not exhibit OPEN-ENDEDNESS. If a different response is taken, perhaps that will turn (22) into a cp-generalization. But for now, I will assume that my account is inapplicable to statements involving probabilistic causation.

6.5 OTHER CHARACTERIZING SENTENCES

When I introduced the semantics for characterizing sentences I accept, I said that some characterizing sentences are not cp-generalizations. I now want to expand on this point, because it allows me to highlight another use of characteristic processes: by appealing to characteristic processes, I can predict which characteristic sentences are cp-generalizations and which ones are not.

Specifically, I predict that characterizing sentences for which no characteristic process is in the offing are not cp-generalizations. Here's why. All characterizing sentences satisfy EXCEPTIONS—that is why, in general, they do not entail the corresponding universal generalization.[37] So cp-generalizations will be distinguished from other characterizing sentences by their OPEN-ENDEDNESS. But when a characterizing sentence exhibits OPEN-ENDEDNESS, it must be possible

[37]The entailment at issue here is semantic entailment, rather than metaphysical entailment. To see what is at issue, consider 'Lions are mammals'. As a matter of the metaphysics of species, lions cannot be anything other than mammals. However, it does not seem that it is part of the meaning of 'lions are mammals' that all lions necessarily are. As far as meaning is concerned, 'lions are mammals' and 'lions have four legs' are on a par.

to give some unifying account of the things mentioned in the open-ended unless-clause. On my view, that is the job of the characteristic processes. When no characteristic process is available, we predict that the generalization does not exhibit OPEN-ENDEDNESS.

This prediction is borne out. Consider again the generalization about English verbs, repeated here:

(23) English verbs form their simple past by adding '-ed'.

Generatively—albeit not etymologically—the reason 'walk' forms its past tense by adding '-ed' is that it is not covered by any of the more specific rules governing exceptions. Clearly, there is no process in the offing, characteristic or otherwise. And indeed, (23) does not exhibit OPEN-ENDEDNESS. The irregular verbs form a finite, closed list.

I just contrasted generative and etymological explanations. Let's expand on the contrast slightly. For a speaker, what makes a verb regular is that the she does not have an explicit lexical entry for the past tense of the verb. The past tense is therefore formed by a default rule that is applied in the absence of more specific rules. By contrast, irregular verbs are represented twice. We learn 'go', and then have a separate entry for 'went', listing it as the past tense of 'go'. However, there are usually historical reasons for why a certain verb has an explicitly listed past tense form. In many cases, they are the result of formerly regular forms that have fallen out of favor. That is why irregular verbs often cluster together, such as (*bring-brought, seek-sought* or *sing-sung, cling-clung*). But at a given time, what makes a verb irregular is that it is listed explicitly as such in a speaker's lexicon. The same theory applies to plurals.[38]

7 Conclusion

I have argued that cp-generalizations which overtly look like generalizations are, in fact, claims about the characteristicness of certain processes. This thesis enables us to respond to one of the most pressing issues about cp-generalizations—the question of how they can be confirmed. Presented most abstractly, I hope to have shown that there is no *special* problem about cp-generalizations. Rather, they are an instance of the much more common phenomenon of theory-laden categorization: some categorizations are good, some are bad, and it is a hard problem to say what makes them so. But that should not lead us to discard our schemes of categorization. In the special sciences, these schemes are reflected by cp-generalizations. So much the better for cp-generalizations.

[38]For a lengthy and readable exposition of this theory, see Pinker (2000). Compare also Haspelmath (2002).

Paper 2

Theoretical Truth and The Presence of Models

1 Introduction

It is the job of the philosopher of science to describe central features of scientific practice, especially ones that have a bearing on perennial philosophical concerns. The use of idealizations in science is one such feature, because idealizations are essentially involved in how theories are confirmed and applied to situations. Theories are neat, the world is not, and idealizations bridge the gap. We see the same strategy in philosophy, for example in formal theories in epistemology, where neat theories such as Bayesianism need to be related to our actual mental states, which are anything but.

Many philosophers of science have suggested that bridging the gap with idealizations comes at a price: truth is not the proper evaluative notion for theories. Theories may be evaluable for truth or not, but regardless of whether they are, their truth is not what matters, and our best account of the role idealizations play in science forces this view on us. I shall challenge this claim by proposing an account of idealizations that is motivated independently of the debate about truth and that makes literal truth an important feature of scientific theories.

Here is the target line of thought. It is incontrovertible that many laws of science, by themselves, make no predictions about the world whatsoever. Newton's Laws, for example, say nothing about what happens in the world. In order to derive predictions, we need auxiliary hypotheses, claims to the effect that things exist, that they're distributed in space in some way, and so on. Very often, these auxiliary hypotheses are false of the actual world. They embody idealizations. Perhaps we cannot measure what the world is like well enough to have a true auxiliary hypothesis. Perhaps considering the false hypothesis makes our lives easier.

Such appeals to idealizations are legion. Indeed, one might think that theories are always applied via some form of idealization or other. One might think that for this reason, the only use for theories is guidance in model building. But nobody expects such models to faithfully replicate the actual world in all respects—indeed, the point of introducing the models in the first place is that they simplify the messy world we find ourselves in. One might draw the further conclusion that the point of theories is at best to describe the models that we use, and hence not to accurately describe the world.

Ronald Giere is one prominent proponent of this view.

> The way scientists practice mechanics, it seems that the literal, exact truth of the laws of motion is irrelevant.[1]
>
> We can [. . .] forget about truth and focus on the details of the similarity [between model and world]. A "theory of truth" is not a prerequisite for an adequate theory of science.[2]
>
> The relationship that does the heavy representational work is not one of truth between a linguistic entity [such as a sentence or proposition] and a real object, but of similarity between two objects, one abstract [i.e., the model] and one real.[3]

Let's give the position on the interpretation of theories that Giere endorses here a name.

[1] Giere (1988, p. 77).
[2] Giere (1988, p. 81).
[3] Giere (1988, p. 82). Similar views are expressed in his (1999, p. 90) and (2004).

[NO TRUTH (NT)] Whatever makes a theory good is not whether it is true or false.

NT is silent on whether theories can be evaluated for truth. According to NT truth, even if available, just does not matter.

Giere endorses a reason for accepting NT. He hints at it in the first passage I quoted: he says that the way scientists practice mechanics suggests NT as the best account of that practice. At greater length, Giere says.

> Those scientists who have taken it for granted that the laws of nature are well confirmed, general statements have obviously not been ignorant of the fact that scientists regularly use approximations. But they have taken this to be a relatively inconsequential fact about science. They have regarded the fact as a matter of only practical, not theoretical, importance. This attitude has been reinforced, I think, by taking logical systems and mathematical theories as the model for scientific theories. In logic and mathematics, at least before the age of computers, approximate methods served only the practical end of making calculations possible, or at least easier. Empirical science is different. Idealization and approximation are of its essence. An adequate theory of science must reflect this fact in its most basic concepts.[4]

[4]Giere (1988, p. 78).

> I suggest calling the idealized systems discussed in mechanics texts "theoretical models," or, if the context is clear, simply "models".[5]
>
> The relationship between some (suitably interpreted) equations and their corresponding model may be described as one of characterization, or even definition. We may even appropriately speak here of "truth." The interpreted equations are *true of* the corresponding model.[6]

As Giere says, idealizations are of the essence in science, and the best account of them incorporates the NO TRUTH claim NT. In Giere's case, that takes the form of saying that theories are true of models. And as the previous quotes from Giere show, the theories are *not* also true of the world: "we can forget about truth."

Giere's view is quite popular.

> It is models rather than abstract theory that represent and explain the behaviour of physical systems.[7]
>
> A scientific theory is basically an idealizational construction. [...] It is [not] a body of factual hypotheses referring to unobservables but still pertaining to actual objects.[8]

[5]Giere (1988, p. 79).
[6]Giere (1988, p. 79).
[7]Morrison (1999, p. 39).
[8]Nowak and Nowak (1998, p. 36).

> theories in physics do not generally represent what happens in the world; only models represent in this way, and the models that do so are not already part of any theory.[9]

Proponents of this view differ on whether theories are actually *about* such idealized models, or whether the relationship between model and theory is more indirect. But they all agree on NT, and they agree that NT is forced on us by the pervasive presence of models in scientific practice.

If NT held up, it would be an extremely surprising result. It would show that the product of our epistemically most responsible activity—organized scientific research—is not properly assessed by asking whether it is true. This rather radical conclusion is unwarranted, as I argue. In the process, I also hope to achieve a clearer understanding of what idealizations are. I do so, because we need to have some good account of how models and idealizations fit into our best scientific practice, their intellectual home. With that in hand, we will be in a better position to assess idealizations in philosophy.

My discussion will also give us a clearer view of an interesting argument against realism in science, what is sometimes called the Many Models argument for anti-realism. It is given by Cartwright (1983) and Morrison (1990). The argument suggests that the practice of using incompatible models and idealizations should incline us towards an anti-realist stance on the theories that license these

[9] Cartwright (1999b, p. 242). See also French and Ladyman (1998), Suárez (1999), and Teller (2001), which has lots of further citations.

models. However, the notion of incompatibility required by the Many Models argument depends on one's view of the relation between theory and models. On the view I propose, the argument fails.

My account in outline is this. Idealizations are epistemic tools. By that I mean that we appeal to them in order to justify our beliefs about what the world is like. Whenever we make use of an idealization, such a use is licensed by a theory we accept. And crucially for my argument against NT, if that theory is false, we are open to criticism. In that respect, the truth of the theories we accept matters, and that is incompatible with NT. The plan of the paper: after some ground-clearing (section 2), I turn to the heart of the paper, the account of idealizations (section 3). With the account in place, I'll discuss an application, the many-models argument against realism (section 4).

2 Two Kinds of Models

I have so far only spoken of models like frictionless planes and two-body idealizations. That is one notion of model. There is another notion of model that comes into the philosophy of science from a different source: mathematical model theory. I want to explicitly distinguish these two kinds of models, because if they are identified, NT follows straightforwardly. In this section, I first present the second kind of model and then show how identifying the two kinds of models leads to

NT.

The mathematical notion of a model comes into the philosophy of science motivated by the short-comings of the logical empiricist analysis of theories, famously called the "received view" by Suppe (1977).[10] Logical empiricists individuated theories syntactically. They had philosophical reasons for this. Chief among them their desire to reconstruct many scientific concepts, such as explanation and confirmation, by relying heavily on the notion of logical consequence, defined syntactically. So they needed an account of theories on which theories entered into syntactic relations. Simply identifying theories syntactically is an extremely natural position if that's the goal.

But identifying theories syntactically also brings with it certain costs. What are intuitively two ways of stating the same theory, say in English and German, turn out to be different theories. They are perhaps closely related, but on the logical empiricist view, they are different theories. Also, whether theories disagree seems to be illegitimately hostage to the vicissitudes of formulation. On the empiricist view, two theories disagree just in case one can derive a contradiction from their conjunction. And that depends on the particular vocabulary employed.

These costs suggest that perhaps theories are not a certain bit of linguistic material, but rather are expressed by that material. So we should identify a theory with "what it says about the world." That notion in turn seems to be well captured

[10] He credits Putnam (1973) with coining the term.

formally by considering the set of worlds compatible with what the theory says, the set of worlds in which it is true.

In mathematics, we have a very similar notion. We begin with a set of sentences, picked out syntactically. We then interpret them: we assign objects to the singular terms, sets to the predicates, and so on. Any such interpretation on which all of the sentences come out true is called a *model* of that set of sentences.[11] The mathematical notion of a model is precise, broad, and metaphysically quite uncommitted. It seems natural to extend this notion of 'model' to refer just to the worlds that are compatible with a theory, since these are the worlds that, if actual, make the theory true.[12] If we then identify a theory with the set of its models, theories aren't hog-tied to the particular form in which they are expressed. Rather, theories are identified with what they say. This is the so-called *Semantic View* of theories.

So a mathematical model of a theory is what the theory is about. Such a mathematical model could be a possible world or some abstract set-theoretic predicate. So if we identify mathematical models with physics-models, we obtain directly that a theory is about these physical models, such as frictionless planes and perfect pendula in the case of mechanics.

But it is important to see that identifying physics models with mathematical models is a substantive move. It is *not* compelled simply by adopting the position

[11] See Boolos and Jeffrey (1989) for an introduction to the mechanics.
[12] This extension is already exhibited at Suppe (1977, p. 96).

that we should individuate theories by what they say. Both the proponent and the opponent of NT can agree with that. They disagree over what it is that a theory says. The proponent of NT says that theories are about physics models; the opponent of NT says that theories are about the world.

I am thus disagreeing with proponents of the semantic view who sometimes suggest that the mathematical notion of a model is not very far removed from the usage of the term "model" in the sciences.[13] I am arguing that, at least conceptually, we need to sharply distinguish between the two.

The point I have belaboured in this section may seem obvious. However, many theorists take as an objection to the semantic view of theories that it does not provide, nor is immediately extensible to, an account of the way practicing scientists use models like a two-body idealization.[14] But that objection *presupposes* that the semantic view is committed to identifying the two kinds of models. Otherwise, there would be no reason to expect the semantic view to be particularly well-suited to give an account of scientific modelling, and hence it would not be an objection to the view that it is not naturally extended into such an account.

[13]See, for example, Suppes (1967) and Van Fraassen (1980, p. 44).
[14]See, for example, Cartwright et al. (1995), Morrison (1999) and Suárez (1999).

3 The Account: Idealizations are Epistemic Tools

I'll begin with what should be a recognizably paradigmatic use of a model, and then argue for a particular interpretation of the episode I present. That interpretation will motivate my account.

Suppose you want to know how fast soap-boxes will go at the bottom of a hill you're considering for a derby. Instead of measuring their top speed, you want to figure it out by computation. To do so, you appeal to an idealization. You consider a block sliding down a frictionless plane, and choose appropriate values for the different variables—weight, incline, etc.[15] You conclude that the block goes 25mph at the foot of the frictionless plane, and that the soap-boxes will reach about that speed, as well.

Here is the interpretation I will defend in this paper. There are three features of the example. We begin with a question: how fast will the soap-boxes go at the foot of the hill? In the course of answering it, we ask and answer another question, this one about the idealization: how fast will a block go at the foot of such-and-such a frictionless plane? We then take the answer to the second question as close enough to the true answer to the first question. In the example, we conclude from the fact that the block goes 25mph that the soap-boxes will, as well. And when all

[15]Throughout this paper, I will consider the idealization of the frictionless plane. In this idealization, there is no friction between the object sliding down the plane and the plane itself. But there is more: the object does not experience any deformation, there is no air-resistance, and a host of other factors obtain, as well. I will tacitly assume throughout that frictionless planes are like that.

goes well, the appeal to the idealization *justifies* us in accepting the answer to the question about the soap-boxes.

Thus quite generally. We begin with an original question about an original situation. Using an idealization to answer it involves three things: answering a question about an idealization, taking that answer to be close to the true answer to the original question, and relying on the consideration of the idealization for evidence for taking the two answers to be close to each other.

For the purposes of this paper, I'll talk of idealizations as *possible situations*. Some regimentation here is required, mostly for clarity's sake, since there is very little agreement in the literature on what idealizations are. There are two important alternatives I rule out. The first is to define idealizations linguistically. The linguistic means by which the idealization is described seems inessential to whatever worries we might have about it. Sometimes, it is also unclear just what idealization one has in mind when one uses a particular form of words, for example, when one describes an idealization metaphorically.[16] I also do not accept a usage on which it is essential to an idealization that it *could not* obtain compatibly with the laws.[17] The frictionless plane certainly has this feature, since there could not be any. But it seems to me that when we appeal to a laboratory result under controlled conditions in order to find something out about the world outside of

[16]The following writers all individuate idealizations linguistically: Barr (1971, 1974), Hempel (1965c), Laymon (1980, 1982, 1989), Rappaport (1998), Rudner (1966), Schwartz (1978).

[17]The following writers take idealizations to essentially be ruled out by the laws: Barr (1971, 1974) and Schwartz (1978).

the laboratory, we have an instance of the same pattern as when we appeal to a frictionless plane. In a specific case, that might be building a wind-tunnel to study how a wing will do in the air or a water-tunnel to study the non-slip condition of fluids.[18]

Proponents of NT are perhaps willing to agree with my interpretation. But I'll argue now that when we rely on an appeal to an idealization to justify us in accepting some answer as close to true, we also rely on the literal truth of a theory. To say that we "rely" on the literal truth of the theory is to say that our epistemic practice is open to criticism when the theory is not true. And that means that NT is not forced on us by an account of the use of idealizations in practice. So what does it take for an appeal to an idealization to play this evidential role? There is an obvious answer: the appeal to the idealization can play the evidential role just in case the subject has reason to think that the answer to the question about the idealization is close to the answer to the original question. While that's uninformative, it does help structure the debate.

A more informative answer will consist of two parts. First, we might hope that there is a kind of reason that is distinctive of idealizations that scientific practitioners can have for believing the answers to two question to be close. Obviously, reasons for belief can come from any source. But in the usual course of scientific practice, it may well be the case that its practitioners have a particular

[18]This last example is taken from Morrison and Morgan (1999, p. 26).

kind of reason for accepting an answer, one that is only available for idealizations. As a hypothesis, this distinctive kind of reason has to do with a feature of idealizations—for example, that the idealization only differs from the original situation in certain ways. This part of the answer is the topic of this paper.

An informative answer also needs to tell us what epistemic attitude someone who uses the idealization needs to bear to the proposition that the idealization has the feature in question. There are many candidates. Perhaps the practitioner needs to know that the idealization has this feature; perhaps she has to justifiedly believe that it does; perhaps she merely has to be in an environment in which experts are about who know that the idealization has this feature. In this paper, I will not take a stand on this issue. For simplicity, I will assume that the relation is knowledge.

There is another issue I will not discuss at all. I'll just assume that we *know* the answer to the question about the idealization. In the usual case, that is just a problem of mathematics, and things can go wrong when we do math. But I'll assume that they don't. So the question for the rest of the paper is what a subject has to know about the idealization in order for an appeal to that idealization to play the evidential role, given that she knows the answer to the question about the idealization.

The rather pedestrian gloss of idealizations as "differing in ways that don't matter" suggests the account I want to give. Roughly, it says that the subject has to know that the causal factors present in the idealization are sufficiently close to the causal factors in the original situation so as not to affect the answer materially. This idea is subject to objections and in need of clarification, and I'll need to do a fair amount of work to flesh it out. That is the task that will occupy the remainder of this section.

3.1 Closeness to the Truth?

One might worry that the train runs off the tracks at the very start: I have been playing fast and loose with the truth, or what is worse, closeness to the truth. I said that answers to the question about the original situation and about the idealization have to be close to each other. Since we are interested in the *true* answer to the original question, that means that I need to make sense of the idea that the claim true of the idealization is close to true of the original situation as well. And to say that closeness to the truth is a contested notion is to put the matter too mildly.

Let me put one distraction to the side first. Being close is a variable notion. One can be closer to something or further away, and what counts as close, *simpliciter*, depends on many factors. But that is not the gravest concern. It is true that the answer to the two questions—about the original situation and the idealization—have to be close enough to each other. Close enough for what? Close enough for the purposes at hand, whatever those may be.

The most pressing problem arises when no clear metric is available with respect to which we can evaluate the "distance" of two things. When we measure the distance of two things in space, space provides a metric. When we measure the distance of two things in time, time does. When we measure the distance between velocities, speeds do. The main problem at issue in the debate about closeness to the truth—verisimilitude—is to define a notion of closeness to the truth when there is no obvious metric to impose. The most prominent case is closeness to the truth for whole theories. The case is interesting for scientific realists who maintain that, even though we may not have true theories, we do manage to get theories that are ever closer to the truth.[19] However, how to say more precisely what it takes for one theory to be closer to the truth than another is very hard.[20]

Given this state of affairs, any account that needs to make sense of closeness to the truth in order to make sense is likely more trouble than it is worth. But I want to make a virtue of the way the problems about verisimilitude arise. Specifically, I'll argue that we should distinguish between two kinds of questions, and depending on the kind of question, the relation between the two answers is different. And importantly for the purposes of my argument, I can motivate the distinction between the two kinds of questions on independent grounds.

[19]This is the strategy followed, for example, by Kitcher (1993) and Weston (1992).

[20]The literature on the question seems to begin with Popper's definition (Popper (1963)), and is quite active now (see, e.g., the surveys in Niiniluoto (1998) and Oddie (Fall 2001)).

3.2 QQ-QUESTIONS

Kinds of questions are picked out by the answers they receive.[21] The questions I have mentioned so far, such as 'how fast will the soap-boxes go?' only form a special case. I'll call them qualitative/quantitative questions, *QQ* for short.

[QQ-QUESTION] QQ-questions are all and only the questions whose answers are quantitative or qualitative claims.

Some examples of QQ-questions:

- How fast are the cars going?
- Where will the Moon appear tonight?
- Will the price of coffee rise in the next three weeks?

This last question counts as having a qualitative claim as answer. In general, claims to the effect that a certain quantity will change in a certain way, but where the claim does not involve by how much the quantity will change, are qualitative claims. For example, questions whose answers are that some quantity will go up or down are QQ-questions with qualitative answers.[22] When a question is a

[21] This isn't as *ad hoc* as one might think. In the formal study of questions in linguistics, the most widely endorsed working hypothesis identifies questions with the set of their answers. See for example Groenendijk and Stokhof (1997).

[22] The sense of qualitative at issue here is the sense one learns in physics class when one is told that one should always get a qualitative answer to any question before starting to grind through the math. Qualitative has nothing to do with "purely qualitative predicates" or any of the usual uses in philosophy.

yes/no question, the answer is the proposition we would follow up the "yes" or "no" answer with, e.g., "yes, the price of coffee will rise in the next three weeks."

These are the questions where the answer to the question about the idealization needs to be close to the true answer to the original question. But that is unremarkable so far. For quantitative answers, the quantity we are concerned with provides the metric for evaluating closeness. For qualitative answers, being close to true is just being true.

3.3 Causal Questions

The much more problematic questions are about causation, "causal questions" for short.

[Causal Question] Causal questions are all and only the questions whose answers are causal claims or claims about causation.

Some examples of causal questions:

- Why did the chicken cross the road?
- How come water expands when frozen?
- What is the significance of air resistance in the soap-box derby?

In order to motivate my treatment of idealizations in causal questions, let me consider an example.

> When a ball is being thrown from one person to another on flat ground, a practiced thrower knows instinctively the right angle at which to throw. He has found out by experience how to throw the longest distance, without having to solve any equations. Again, neglecting air resistance, we can make an estimate of how steeply he should in fact throw. We have already found that, if something is thrown vertically upwards with an initial speed v_0, it takes a time $2\frac{v_0}{g}$ to reach the ground again, since it has a constant downward acceleration g, and returns with the same speed v_0 downwards. If the ball is now thrown at an angle θ to the horizontal, it is in the air for a time $\frac{2 \cdot v_0 \cdot \sin\theta}{g}$. Since its horizontal speed is constant in the absence of air resistance and has the value $v_0 \cdot \cos\theta$, in this time the ball travels a horizontal distance
>
> $$\frac{2 \cdot v_0^2 \cdot \cos\theta \cdot \sin\theta}{g} = \frac{v_0^2 \cdot \sin 2\theta}{g}$$
>
> The largest possible value of $\sin 2\theta$ is one, when $\theta = 45°$. If air resistance is negligible, $45°$ is therefore the best angle at which to throw.[23]

This passage might be presented as an answer to the question: at which angle will a projectile fly farthest, holding its muzzle velocity constant? But it might also be

[23] Grant and Phillips (2001, pp. 42-3).

offered in response to the question: "Why is 45° the angle at which a projectile flies farthest?" Let's focus on this latter use. What role does the appeal to the idealization play here?

In a perfectly good sense, the idealization is not part of the answer to the question. To see this, not that (1) sounds bizarre.

(1) The best angle for someone to throw a ball is 45° because in a frictionless medium, that's the angle that maximizes the distance of the throw (holding constant initial velocity).

And any sentence like (1) that follows 'because' with some claim directly about the idealization will be equally bizarre. Instead, the answer to the question is (2).

(2) The best angle for someone to throw a ball is 45° because in that situation, the distance of the throw depends on the launch angle (holding constant initial velocity), and 45° maximizes that distance.

The answers to the question about the original situation and the idealization aren't just similar, they're the same. The reason that the projectile flies farthest in the idealization when launched at 45° is the very same reason that the ball flies farthest when thrown at that angle in the original situation.

One might think that the idealization and the original situation differ, and so that they do not exhibit the same causal factor. After all, in the one situation, a ball is thrown, in the other, a projectile is launched. But I think this is a mistake. Quite generally, some features of a situation are causally relevant, some features are not. For example, when one billiard ball collides with another and because of that, the other moves, some features of the first billiard ball are causally relevant to the impact, others are not. The first ball's momentum belongs in the first class, its color in the second. In the example of the projectile, the shape of the projectile is not a causally relevant feature, nor is the fact that it was thrown as opposed to being launched by some other means.

That means that, for causal question, we have the following structure: when a subject appeals to an idealization, she does so for evidence that the answer to the causal question about the idealization *is the same* as the answer to the causal question about the original situation. In my argument for this claim, I have only considered one kind of causal question, the kind that receives an answer of the form *p because q*. But it should appear that similar considerations are available for other causal questions, as well. The distinction between causal and QQ-questions allows me to respond to the worry about verisimilitude. The worry arises when we try to compare propositions about matters of fact for which no obvious metric presents itself, or where several different metrics would yield different results. The distinction between QQ-questions and causal questions disarms these worries. Causal questions are the questions for which the worry about closeness is most acute. The proper evaluation for using idealizations to answer causal questions is not closeness to the truth, it's being the same. On the other hand, in the case of QQ-questions, there is an obvious metric. So my account of idealizations depends on closeness to the truth only where that notion is uncontroversial.

The question now becomes what it takes for such an appeal to be successful, i.e., what it takes for the appeal to actually provide the justification. I'll begin with causal questions. As I said, the platitudinous gloss that good idealizations only differ from the original situation in ways that don't matter to the purposes at hand suggests the account. The platitude has two aspects to it. First, the idealization takes account of all of the causal factors that are present in the original situation. The objects are duplicated, as are the forces. And second, the idealization ignores some of these factors, ones that, even though ignored, are still close to what they are in the original situation. Both steps are important, although it might seem that the first step is superfluous. After all, the result of "taking friction into account" and then setting it to 0 seems to be the same as simply not taking friction into account. But epistemically, there clearly is a difference between thinking about friction and seeing that one can disregard it, and not thinking about friction at all. It seems plausible that the appeal to the idealization can play its epistemic role only if one has consciously ignored friction, rather than forgotten it completely.

So the idealization has to take into account all of the causal factors that are present in the original situation, and the causal factors in the two situations have to either be the same, or they have to be close to each other. One important fact about causal factors for my purposes is that what they are depends on what else is true of the situation. So long as we consider motion in a frictionless medium, the shape of the projectile does not matter, and hence is not a causal factor. On the other hand, once we consider air resistance, shape becomes a causally relevant factor, at least to the extent that it influences the amount of friction the object experiences.

Appealing to causal factors that are close without being the same may raise the concern of verisimilitude all over again. Very similar considerations apply here. Some causal factors are the same in the idealization and the original situation, so there is no problem of verisimilitude for them. The problem only arises when the causal factors are different, and thus have to be close to each other. But for many of these factors, the great advance of science is precisely to provide metrics that allow for the kind of evaluation of closeness that is required. I cannot give a general argument that metrics are provided for all causal factors that have to be close. But a survey of the factors involved in the case of the projectile certainly suggests as much: we can measure friction with friction-coefficients and we can measure the curvature of the Earth's gravitational field. Thus my account is a hostage to fortune: it may be refuted if we can show that there are causal factors that would have to be close, but for which the theory the subject accepts does not provide an obvious metric.

Recall now the shape of the proposal: I said that I am interested in what the subject has to know about the relationship between the idealization and the original situation in order for the idealization to play the evidential role. Right now, my account says that one is justified in accepting an answer by an appeal to an idealization only if one *knows* what the causal factors in the original situation are, and how they interact. That seems too strong, since it entails that nobody with a false theory about the original situation can have justified beliefs on the basis of appeals to idealizations. And that includes just about everybody. It seems better to require merely that the idealization and the original situation share the same causal factors according to a theory the subject reasonably accepts. That way, a Ptolemaic astronomer can justifiably rely on her Ptolemaic models, even though these models get the causal factors wrong.

We now have the first part of my theory of idealizations. Let's summarize it.

[EVIDENCE FOR CAUSAL QUESTIONS]

> An appeal to an idealization I can provide evidence to a subject for accepting an answer to a further question about I as the answer to a causal question about the original situation if and only if
>
> the subject knows that
>
> - All and only the relevant causal factors in the original situation (according to a theory the subject reasonably accepts) are taken account

of in the idealization (where being assigned 0 counts as being taken account of.)

- All the factors in the idealization are either the same as they are (according to that theory) in the original situation, or they are close enough to what they are (according to that theory) in the original situation.

This collects the points I have made over the course of the last few paragraphs. The answers to the question about the original situation and the idealization have to be the same; the idealization has to take account of all the causal factors in the original situation; all the causal factors are either the same or close enough; and finally, what matters about the original situation is what the subject reasonably believes to be true about it, though perhaps not what in fact is true.

Let us say that when according to a theory, the idealization and the original situation have the proper relation—all the causal factors are either identical, or close to each other—the theory *licenses* the idealization for that question. Thus, what the subject has to know in order for the idealization to play the evidential role is that the theory licenses the idealization for that question.

3.4 Arguing against NT

So far, I have only discussed using idealizations to answer causal questions. I will turn to QQ-questions in a moment, but since their treatment is slightly more involved, I want to show first how this account allows me to mount an argument against NT. The argument will go essentially the same for QQ-questions, but the extra layer of complexity there makes it a little bit harder to see.

I said at the beginning that the way to argue against NT was to present an account of idealizations according to which we are committed to the truth of a scientific theory. Commitment here is a matter of being open to criticism in our reliance on the idealization if the theory is false.

The account of causal questions bears this idea out. In order for an idealization to play the evidential role, the subject using the idealization has to know that, according to her theory of the original situation, either the idealization recognizes the same causally relevant factors as are present in the original situation, or that the factors are close to each other. That in turn requires having a theory about what the causally relevant factors in the original situation are. And *that* theory includes what we usually call a scientific theory, although it requires much more besides, to wit, what the more specific facts of the situation are. But no matter. That more is required is immaterial. To see how this works in the specific case of the projectile I have been discussing, the subject has to know that according to her theory, in the actual world a launched object is subject to gravity, and she has to know how much friction one can safely ignore according to that theory. That, in turn, requires knowledge of how according to the theory, the forces present interact to yield an overall acceleration. But that is just the Newtonian theory that according to proponents of NT is not properly evaluated for truth.

So my argument against NT is this: if a subject's appeal to an idealization is licensed by a false theory of the original situation, she is open to criticism. For that reason, the truth of the licensing theory matters. It's important to my argument that commitment to the truth of the licensing theory is spelled out as being open to criticism, rather than not being justified, if the licensing theory is false. Even when the licensing theory is false, a subject may still justifiably believe that the answer to the question about the idealization is close to the answer to the original question. Indeed, the subject may even know the proposition she accepts on the basis of an appeal to an idealization that is licensed by a false theory. It seems plausible, for example, that even a Ptolemaic astronomer knew where in the sky the moon would appear a day or two in the future, even though their theory was mistaken. But even in that case, she is still open to criticism. Consider how we would react if we learned that our own licensing theory is false. Once we learn this, we cannot appeal to the idealization licensed by that theory anymore, at least not without further ado. That further ado will take center stage in my account of QQ-questions, to which I now turn.

3.5 QQ-Questions

The initial example about the soap-boxes with which I began suggests that the account of idealizations for QQ-questions will parallel that for causal questions quite closely. Unfortunately, it does not. But the parallel is an important starting point. Recall that in the example of the soap-box, we considered a frictionless plane that had roughly the same incline as the hill the soap-boxes were to go down, and the soap-box corresponded to an object sliding down that plane.

Indeed, *one* way for the appeal to an idealization to provide evidence for accepting an answer to a QQ-question is for the idealization and the original situation to match as they do for answering causal question. So this criterion provides a sufficient condition.

[EVIDENCE FOR QQ-QUESTIONS (FOR NOW)]

> An appeal to an idealization *I* can provide evidence to a subject for accepting an answer as close to the true answer to a QQ-question about an original situation if
>
> the subject knows that
>
> - all and only the relevant causal factors in the original situation (according to a theory the subject accepts) are taken account of in the idealization (where being assigned 0 counts as being taken account of.)
> - All the factors in the idealization are either the same as they are (according to that theory) in the original situation, or they are close enough to what they are (according to that theory) in the original situation.

As before, it is crucial that the theory the subject accepts tells her how the causal factors interact, and hence what counts as being close enough for the purposes of answering a particular question about the original situation. Notice that disregarding friction is fine for answering how fast the soap-boxes are at the foot of the hill. But suppose we want to ask how far the soap-boxes roll, once they reach the level surface at the foot of that hill. In that case, we should not appeal to an idealization of a frictionless, level plane over which a block slides with some initial velocity. After all, in that situation, the block experiences no force influencing its velocity in the direction it is going, and hence never stops.

Let me now turn to how the account is incomplete. Consider the following pair of cases. In the first, a physics student studying mechanics is given this problem.

[AIR] A bowling ball rolls down an inclined board near the Earth's surface. The board is 3m long, its angle of incline is 30°, its surface is more or less smooth. The bowling ball weighs 5kg. What is the bowling ball's velocity at the bottom?

In her exam booklet, she describes an idealization in which a block slides down a frictionless plane. The block is assigned the same mass as the ball, the incline of the frictionless plane is 30°, and the force of gravity is assigned close to its actual value. She gets the math right, and she receives full credit.

Now she is given the following problem.

[OIL] A bowling ball rolls down an inclined board near the Earth's surface. The

> whole setup is submerged in heavy motor oil, and the ball is pushed down the plane with a stick with such-and-such force. What is the ball's velocity at the bottom?

Though this is not explicitly stated in the problem, the various forces are such that the velocity at the bottom is the same as in AIR. At this point, there are two ways for the student to go about solving the problem.

She could describe a relatively complex idealization that takes account of the viscosity of the oil and the extra force exerted by the stick. It would still ignore the friction exerted by the plane itself, as well as any deformation the ball might undergo. She would then have to do a fair bit of complex mathematics. If she successfully solves the mathematical problem, she gets full credit. But she could also figure out that all the extra forces cancel each other out, and appeal to the simpler idealization she used to answer the question in AIR. If she displays this knowledge and gets the math right, she gets full credit.

Let's say that when a question about one idealization gets an answer that is close to the answer to a corresponding question about another idealization, the two idealizations are *case-equivalent for that question*. In the example just rehearsed, the simple and the complex idealizations are case-equivalent for the question how fast the ball is rolling at the bottom of the inclined plane. With this bit of terminology in place, here is the second way for an idealization to play the evidential role. The appeal to the idealization can play the evidential role if the subject knows that the idealization is case-equivalent to an idealization that could provide evidence via the first avenue. That leaves us with the final account for QQ-questions.

[EVIDENCE FOR QQ-QUESTIONS (FINAL)]

> An appeal to an idealization I can provide evidence to a subject for accepting an answer as close to the true answer to a QQ-question about an original situation if
>
> the subject knows *either* that
>
> - all and only the relevant causal factors in the original situation (according to a theory the subject accepts) are taken account of in the idealization (where being assigned 0 counts as being taken account of.)
> - All the factors in the idealization are either the same as they are (according to that theory) in the original situation, or they are close enough to what they are (according to that theory) in the original situation.
>
> *or* that the idealization is case-equivalent to one that satisfies the first condition.

Even with case-equivalence, we can mount an argument against NT here. For idealizations that play the evidential role via the first way, the argument against NT is exactly as before: when a subject relies on an idealization licensed by a false theory, she is open to criticism. But notice that the argument does not rely crucially on the fact that the idealization is licensed *directly* by the theory. An appeal to an idealization licensed by a theory by way of being case-equivalent to an idealization licensed directly is no less licensed by the theory for all that.

Let me stress this result, in large part because the notion of case-equivalence serves to cover a wide range of idealizations. The example I used to introduce the notion involved two very similar idealizations. But that need not be the case. Case equivalence also explains why engineers who accepted relativity theory could appeal to Ptolemaic astronomical models to predict the position of the stars in the short term, or why we can appeal to geometrical optics.[24]

Let me now make a cautionary remark. It is a tempting heuristic to assimilate the distinction between QQ- and causal questions to the distinction between prediction and explanation. But the two distinctions do not coincide, and when they come apart, the prediction/explanation distinction falls in the wrong place. It certainly seems possible to predict what the causally relevant features of something will be, and hence answering a question about such features will sometimes be making a prediction. But that question is nonetheless a causal question, not a QQ-question.

[24]The first example comes from Kuhn (1962), the second from Shapere (1969, pp. 152-4).

3.6 Compare and Contrast?

I've said that someone using an idealization needs to believe that the causal factors in the idealization and the original situation line up properly. That suggests that she needs to know what the causal factors in an idealization are, what they are in the original situation, and that she then needs to compare and contrast. That in turn might suggest that her knowledge of the causal factors in the original situation has to be *suitably independent* of her use of idealizations to answer questions about that situation. And this last bit is quite unacceptable, indeed. We often find out what the causal factors in a situation are *by* finding out which idealizations deliver the right results in lots of different situations.

For all I have said, I need not deny that. I'm only making claims about what a subject needs to know, though perhaps other epistemic attitudes are alright here, as well. I say nothing about *how* one might come to know this. The fact that a certain idealization has worked well in many and different cases is precisely the kind of evidence practitioners have for believing certain hypotheses about the causal structure of the phenomena they study. So for example, it is quite compatible with my account that instrumentally successful idealizations guide theory choice: if a model is instrumentally successful—if it allows us to repeatedly make predictions across a wide range of cases that are borne out in practice—we can hypothesize that the causal factors the idealization posits are the causal factors in the original situation. So my account can accommodate the concerns of writers like Morrison and Morgan (1999), Morrison (1999), and Suárez (1999) who emphasize that idealizations or models often play this role.

4 The Many-Models Argument

At the beginning, I said that we can also get a clearer appreciation of the Many Models argument against realism. In this section, I'll do four things. First, I'll reconstruct the argument in valid form. Doing so will make the discussion more manageable (4.1). One of the crucial notions that appear in my reconstruction of the argument is that of two models' being incompatible. I'll argue that if one accepts NT, one notion of incompatibility is appropriate (4.2). However, if one does not accept NT, the argument fails as stated (4.3). After making this objection, I'll present a modified form of the argument that *may* be successful, although the examples usually cited in its support fail to do the job (4.4).

4.1 The Argument

The Many Models argument is given both by Cartwright (1983) and Morrison (1990). Thus Cartwright:

> The problem for realism is the first stage of theory entry. If the models matched up one-to-one, or at least roughly so, with the situations we study, the laws which govern the model could be presumed to apply

> to the real situations as well. [...] Different incompatible models are used for different purposes.[25]

I think the idea behind the argument is roughly this. Idealizations are situations that differ from the original situation in inessential ways. Perhaps they lay bare the causal connections in a way that help us get a grip on complex situations in the actual world. That entails that idealizations are quite similar to the actual world we are studying. Or more precisely, as I have argued above, idealizations are quite similar to the world as we conceive of it through our theory. But suppose that we find scientists using idealizations that are really quite different. Then the world has to resemble things that do not resemble each other at all. If the theory nonetheless tells us that using these disparate idealizations is alright, the theory itself cannot take its conception of the world seriously—we should take an antirealist attitude towards any theory that licenses such disparate idealizations.

Three features of this argument differentiate it from more well-known antirealist fare. For one, it does not presuppose any of the usual distinctions antirealists appeal to, such as the distinction between what is observable and what is not.[26] Arguments that rely on these distinctions start off badly. This one does not.

Second, the argument works theory by theory, rather than trying to establish all at once a conclusion that holds for all of science. To contrast this with a more

[25]Cartwright (1983, p. 158).

[26]Or the distinction between phenomenological and fundamental laws mentioned by Cartwright (1983).

famous example, van Fraassen's anti-realist argument relies on a very general epistemological premise about what beliefs experimental results justify. On his view, experimental results never justify the belief that the unobservable objects and properties posited by the theories really exist.[27] By contrast, the present argument is an argument schema that needs to be instantiated theory by theory. If a particular theory makes the premises true, we should be anti-realists about that theory. If another theory fails to make these premises true, we should not be anti-realists about that one.[28] Potentially, we can have a much more differentiated attitude towards different theories if we endorse this kind of argument than if our anti-realism rested on fully general grounds.[29]

Third, if we adopt anti-realism about a theory on the basis of the present argument, we can have a much more sympathetic attitude towards what scientists actually say. Compare again with global anti-realism accepted on very general epistemic principles. Defenders of anti-realism adopted on these grounds have to say that scientists make some systematic error in their assertions when they say things like "There are electrons in the cloud chamber", no matter what else these

[27] See Van Fraassen (1980, chp. 2).

[28] As Sklar points out, the anti-realist arguments most influential among scientific practitioners do not have as premises completely general claims that could apply to any science (see Sklar, 2000, pp. 11-12). So the form of this argument may well be more compelling than global anti-realist arguments. Kincaid (2000) makes a similar point about the social sciences: realism about this or that science turns out to be a broadly empirical issue.

[29] I take it that something like this complex set of attitudes is what Cartwright (1999a) argues for. If this really is to be an advantage, it must be possible to draw distinctions between theories. On a very strongly holist conception, anti-realism in one bit of the net may spread to anti-realism everywhere. But on most conceptions of theories, the anti-realist conclusions for one theory need not entail anti-realist conclusions for other theories.

scientists do. By contrast, the Many Models argument takes as its starting points the attitudes of scientists towards their theories, as they are exhibited in their use of idealizations licensed by these theories. If they persist in talking realistically about these theories, we at least have reason to say that their practice is internally inconsistent, and thus stand on firmer ground when we reinterpret some of their claims.

So let's cast this argument in valid form. To make the exposition easier, let's say that the features the idealization countenances, such as the objects and forces it posits, and the kinds of interactions between these things, are the idealization's *ontology*.

[THE MANY MODELS ARGUMENT]

(i) Scientists sometimes use idealizations (I and I^*) with incompatible ontologies.

(ii) If scientists use idealizations with incompatible ontologies, they are not committed to both of these ontologies.

∴(iii) Scientists are not committed to the ontologies of both I and I^*. (from (i) and (ii))

(iv) If someone is not committed to the ontology of an idealization, she is not committed to the ontology of the theory that licenses that idealization.

(v) The same theory T licenses I and $I*$.

∴(vi) Scientists are not committed to the ontology of T. (from (iii)-(v))

(vii) Scientists endorse T as successful.

∴(viii) Scientists endorse T and they are not committed to its ontology. (from (vi) and (vii))

Once (viii) is in hand, we can refute realism about the theory T, since to be a realist about T requires one to be committed to its ontology.

4.2 Incompatibility, Assuming NT

The argument relies on a notion of incompatibility between idealizations, which appears in the first two premises. If we accepted the No Truth claim NT, it would not be very hard to find such incompatible idealizations. Suppose we accepted NT in the form of identifying the mathematical models of a given theory T with its physical models, that is, its idealizations. And an idealization counts as an idealization of T if it is licensed by T. So suppose two idealizations are both licensed that disagree on whether a certain force exists. For example, an idealization that posits no friction and an idealization that posits some friction disagree in this way. If we include both of these idealizations among the models of the theory, the theory has straightforwardly incompatible models, and is thus shown to be inconsistent. If we accept it anyway, we do so not because we think that the theory is true, but because it might serve some instrumental purposes.

So if we accept NT, it is very easy to find idealizations that are incompatible in the sense of the Many Models argument. All it takes is for there to be two models that disagree on what there is.

4.3 The Failure of the Argument

Once we reject NT, however, a different notion of incompatibility is required. To show that this is the case, I'll argue that the Many Models argument as I have just presented it fails. Specifically, I want to focus on the following premise:

(iv) If someone is not committed to the ontology of an idealization, she is not committed to the ontology of the theory that licenses that idealization.

That premise fails because an idealization might be licensed via case-equivalence. We might make use of an idealization that we think gets the ontology of the original situation wildly wrong. Nonetheless, our theory may tell us that appealing to that idealization will give us an answer that is close enough to the true answer we were seeking. So the ontology of an idealization need not reflect the ontology of the theory that licenses it when the licensing goes via case-equivalence.

The earlier example I used to motivate case-equivalence will help here. Suppose we want to know where the moon will appear tonight. One idealization we could appeal to is a Newtonian two-body idealization that eliminates all heavenly bodies other than the Earth and the Moon. There will be differential equations to solve, and that will be computationally challenging. Another idealization one could appeal to is a Ptolemaic model, which would be computationally much less challenging. Both idealizations will deliver pretty much the same answer to the question, and so they are case-equivalent for *this* question. But that doesn't mean that we need to be anti-realists about the starry heavens without. That means that one can be uncommitted to the ontology of an idealization while being fully committed to the ontology of the theory that licenses the idealization, because the idealization may have been licensed via case-equivalence.

But the case-equivalence response has its limits. Some questions constrain just how disparate models can be that are case-equivalent for those questions. These are the causal questions. Here's why. An idealization can be used to answer a causal question about some original situation only if the causal factors are the same in the idealization and the original situation, at least at the relevant level of description. So the answer any idealization gives to a causal question will mention the causal factors of the original situation, at the relevant level of description. Any idealization that is case-equivalent with respect to that causal question will mention the same causal factors.

Return to the example of the why 45° is the best angle to throw a ball. We appealed to an idealization, and the answer we got was that when there is no air resistance, the distance is maximized by a launch angle of 45°. Any case-equivalent idealization will give the same answer (that's what makes it case-equivalent). So if we were to use a disparate idealizations to answer a causal question, that would not be licensed by appeal to case-equivalence. The realist would be in trouble. For instance, if we could answer causal questions about the heavens both by appealing to a Newtonian model and a Ptolemaic one, that would be a *prima facie* embarrassment for realism about those heavens.[30] That suggests a reformulation of the Many Models argument.

4.4 The Revised Argument

Right now, the argument relies on the premise that

(iv) If someone is not committed to the ontology of an idealization, she is not committed to the ontology of the theory that licenses that idealization.

And (iv) fails in the complete generality in which it is stated, since it fails for idealizations used to answer QQ-questions. That was the upshot of thinking about

[30] So far, my response to the argument is similar to that in Kukla (1994). He says:

> realists can rationally avail themselves of multiple models even if they think that they already possess the true theory of the phenomena. For example, they may predict an experimental result by using an alternative theory which they believe to be false, so long as they have reason to believe that the false theory's prediction is sufficiently close to the true theory's over a specifiable domain of phenomena. They might choose to do so because the analysis from the true theory is prohibitively complex. Realists who believe in the theory of relativity are not thereby barred from tackling certain classes of problems with the apparatus of classical mechanics.[31]

However, he seems to think that this is the end of the argument. But against Kukla, I argue that the case-equivalence response he and I both seem to favor has its limits. As I argued in the main text, the crucial premise (iv) holds for causal questions.

engineers who believed relativity theory nonetheless using Ptolemaic models of the heavens. They were committed to the theory that licensed the model, but not to that model's ontology.

However, the limits of case-equivalence suggest that (iv) holds in a restricted form, to wit (iv′).

(iv′) If someone is not committed to the ontology of an idealization used *in answering a causal question*, she is not committed to the ontology of the theory that licenses that idealization *for answering a causal question.*

To make the argument valid, we need to make corresponding changes in another premise. We now have (i′) instead of ((i).

(i′) Scientists sometimes use idealizations (I and I^*) with incompatible ontologies *in answering causal questions.*

This reformulation of the Many Models argument raises a question. Do its proponents exhibit cases in which disparate models are used to answer causal questions? Not obviously. Cartwright mentions situations in which we want to compute how one property of amplifiers measured in numerical terms changes as another property changes, as well.[32] That is best interpreted as answering a QQ-question.

A much more promising, but also ultimately unsuccessful, discussion comes from Morrison. She discusses the example of various laws one can derive about

[32] See Cartwright (1983, p. 103-4).

gases. Boyle's Law (aka the ideal gas law) is derived by thinking of the molecules of a gas as billiard balls. Van der Waals equation is derived by thinking about these molecules as weakly attracting rigid spheres. And finally, so-called transport equations receive another model again, point centers of inverse power repulsion.[33] Most charitably, we interpret this example as about explanation: we ask why the various laws hold of gases, at least within certain ranges. The different derivations are then presented as explanations for why they do.

But what Morrison needs in order to make out her argument is not just that the models are incompatible, *simpliciter*. We might say that two models are incompatible just in case they differ from each other, and the models Morrison mentions clearly do. But what matters is whether these models posit different causal factors. To see this point, return to the question about the best launch angle. Two case-equivalent idealizations might be one where the projectile is square and one where it is round. These two idealizations are incompatible. If one of them obtains, the other doesn't. But since there is no air-resistance, the difference does not matter. The causal factor at work is just the presence of a projectile, not its shape, the idealizations thus posit the same causal factor, in this case involving a projectile.

So what about the different models used to derive (and thus, in this case, explain) the different laws describing gases? The two laws Morrison focuses on

[33] See Morrison (1990, p. 319).

are the ideal gas law and the van der Waals equation.[34] The ideal gas law says that in an ideal gas, pressure, temperature, volume, and amount are proportionally related. An ideal gas is one made up of molecules that have no extension and engage in perfectly elastic collisions. The law is then derived on purely mechanical principles. So far, we're only discussing the idealization and what is true of it. In a standard discussion of the ideal gas law, the derivation also helps to explain why real gases are approximately described the by the ideal gas law, at least under certain conditions. The answer is that when the molecules making up the gas are spaced sufficiently far apart, they engage in nearly enough elastic collisions, and the volume they take up is very small relative to the volume the gas takes up.

The reason the molecules need to be "sufficiently far apart" is that actual molecules, as opposed to the bodies making up an ideal gas, exert attractive forces. That means that the pressure one actually observes is less than that predicted by the ideal gas law. In order to take account of the attractive force between molecules, we formulate the van der Waals equation. It takes into account the attractive forces between molecules, as well as the volume that is unavailable for molecular motion.

For our purposes, the crucial fact is that both equations, the ideal gas law and the van der Waals equation, are motivated by the same underlying representation

[34]The exposition here follows the discussion in Grant and Phillips (2001, pp. 422-5, 462-3).

of the molecules making up a real gas.[35] In particular, one appeals to nuclear potentials to explain both why molecules engage in something like elastic collisions (as required for the derivation of the ideal gas law) and why there are attractive forces between molecules (as required for the van der Waals equations). That means that, at the relevant level of description, each of the models involved posit causal factors that are really present in the original situation we are asking a causal question about.

To finish this discussion, here is one last arena in which the many models argument can get traction: the wave-particle duality of light. Light exhibits certain features that suggest that light is a wave, most notably, interference effects. But it also exhibits phenomena that suggest that it is made up of separate objects, photons, such as the photoelectric effect. So perhaps we explain why light exhibits interference phenomena by appealing to a wave-model of light; and perhaps we explain why light exhibits the photoelectric effect by appealing to a particle-model of light. Light can't be both a wave and a particle, so we seem to have a genuine case in which the many models argument gets off the ground.

If the argument really does get off the ground, that is due to a significant thesis about how physics ought to be interpreted. It is not physical common sense. According to one textbook, the wave-model is predictively accurate, but not explanatorily adequate:

[35] Indeed, in my physics textbook, nuclear potentials are introduced before features of gases are derived.

> The property of the photon of behaving as both a particle and a wave is called **wave-particle duality**. Although at first sight it may seem paradoxical, it is not really very mysterious. After all, once it is recognized that matter is made up of atoms, it is only possible for *individual* atoms to be excited by light. The wave and particle aspect are describing different parts of a system's response to light. The wave theory gives advance information of the probability that light will reach different places; the particle picture is no good for this because it cannot account for interference. But once an interaction happens, as for example in a Compton scattering, the scatting is determined by the energy and momentum of a particular photon.[36]

> The wave pattern [reflected by interference phenomena] is carried by each separate photon, which has its own probability distribution independently of other photons.[37]

That is, if we want to *predict* the distribution of light and dark in an interference experiment, we make use of the wave-theory of light. However, when we want to explain *why* there is a certain pattern, we appeal to the model of the photon as particle. As it happens, it's a very strange sort of thing, since it has its own probability distribution. That, by itself raises a host of issues about realism, but

[36]Grant and Phillips (2001, p. 321), emphasis in the original.
[37]Grant and Phillips (2001, p. 321).

they are orthogonal to the many models argument.

The particular examples the proponents of the many-models argument marshal thus fail to establish anti-realist conclusions, since they fail to establish the first premise of the revised argument. But I have not given a general argument that there are no such cases, or even that there could not be. For all I have said, we may yet come across some.

5 Conclusion

I have argued for two claims. Our best account of idealizations does not, by itself, force upon us the claim that theories are not properly evaluated by asking whether they are true. And once we are clear on how idealizations work, an otherwise promising argument for anti-realism fails.

There are uses of idealizations that I have not discussed, since they do not fit into either the mold of causal or QQ-questions. Among them are questions such as "what should I investigate next?" This is a question Hesse (1966) has drawn out attention to. The fact that scientific theories provide guidance not just about what the world is like, but centrally also about what it makes sense to think about next is a point Bromberger has emphasized.[38] That is an important objection to standard logical empiricist construals of theories, since their theory of theories does nothing to help us see how theories can provide this guidance. But for my

[38] See Bromberger (1992a,b).

purposes, the fact that theories do more than just represent what the world is like is immaterial. Likewise, the fact that idealizations are used at times to answer questions about things other than what the world is like is immaterial. I have sought an important class of cases in which idealizations are used, and in which the best account of idealizations argues against NT, not for it.

Paper 3

Can There Be a Theory of Explanation?

1 Introduction

Many explanations are given by offering because-claims, claims of the form *p because q*. The classic discussion in Hempel and Oppenheim (1965) takes these claims as its target of analysis. We wonder why the radiator in the car cracked overnight, and learn that it did because it was full of water and the water froze. Because-claims entail their corresponding conjunction, but are not entailed by it: *p because q* entails *p and q*, but not vice versa. A theory of explanation seeks to tell us what kind of additional information is contained in a because-claim that makes it strictly stronger than the corresponding conjunction.

Hempel had an answer to that question. On his view, *p because q* is true (roughly) iff the sentence $\ulcorner q \urcorner$ is part of a derivation that contains some statements of laws essentially, and whose conclusion is $\ulcorner p \urcorner$. This leaves open just what epistemic relation one needs to bear to the various premises in the deduction, and the view only covers some cases, at any rate, since it does not cover explanations of laws. That is the deductive-nomological (DN) account of explanation.[1]

The DN account of explanation has been opposed, the opposition taking many different forms. There are causal theorists of explanation, such as Lewis (1986) and Salmon (1984), who suggest that a because-claim asserts a causal connection between the fact that p and the fact that q.[2] There are unificationists such as Friedman (1974) and Kitcher (1989, 1991), who suggest that because-claims assert that appeal to the fact that q allows us to fit the fact that p into an epistemic pattern along with lots of other facts. While these theorists disagree on the true philosophical account of explanation, they agree that it is possible to give such an account. They share the presupposition that there is a uniform phenomenon of explanation, and that one can give a philosophically illuminating account of the

[1]The initial statement of the DN account is given in Hempel (1965b), and elaborated in Hempel and Oppenheim (1965) and (1965a).

[2]They are certainly not the only ones. Indeed, it seems fair to say that causal theories of explanation are dominant right now. See, e.g., Strevens (2004) and Woodward (2003), as well as the theorists cited in note 20.

phenomenon. Let me mark this commitment.

[UNIFORMITY] There is a single phenomenon of explanation that we can give a philosophically illuminating account of.

All of the theorists who endorse the uniformity assumption accept that there is a distinctive kind of explanatory information that is contained in a because-claim. Some will say it is information about nomological subsumption, others about causation, others about unification. But they all agree that we can give an illuminating answer to the question: what further information is contained in a because-claim that is not already contained in the corresponding conjunction?

Attitudes towards UNIFORMITY come in degrees. At one end of the spectrum are the theorists I've already mentioned: covering law, causal, and unificationist. They claim that we can go very far in giving a philosophical account of explanation. Somewhat removed from this end of the spectrum is the account of Miller (1987). On his view, we can give a necessary condition on something's being an explanation: information is explanatory *only if* it is causal. However, this condition is not sufficient. He says that information is explanatory if and only if it is *adequately specified* causal information, and the criteria for adequacy depend on the particular science we are pursuing, as well as the stage of that science's development. On such a view, a significant part of the project of giving an explanation is a posteriori and field-specific. Philosophy leaves that issue unaddressed.

Van Fraassen (1980) has issued the most radical challenge to UNIFORMITY.

> [scientific explanation] is a use of science to satisfy certain of our desires; and these desires are quite specific in a specific context, but they are always desires for descriptive information [...] The exact content of the desire, and the evaluation of how well it is satisfied, varies from context to context. *It is not a single desire, the same in all cases, for a very special sort of thing, but rather, in each case, a different desire for something of a quite familiar sort.*[3]

He suggests that what additional information is asserted when we make a because-claim depends on the context. And he makes the further claim that there are hardly any constraints on the kind of further information the context could thus determine. In principle, any further relation between the fact that p and the fact that q could be such that, if it holds, *p because q* is true. In some situations, we may care about causal information, in others about unifying information, and in others still information about which events temporally preceded others. Certainly, on a particular occasion, very specific information will be required in order for a because-claim to be true with respect to that context. But the interests that determine which information that is are not interests in having an *explanation*. They are specific interests determined by the background beliefs of the conversational participants that go beyond their desire for an explanation *per se*.

If this view of explanation is right, the project of giving a philosophically il-

[3] Van Fraassen (1980, p.156), my emphasis.

luminating account fails. We can of course say that a because-claim is true in virtue of satisfying whatever constraints are imposed by the context. But we cannot say what these constraints might be. Potentially, anything could play this role. Philosophically, we cannot move beyond the platitude that explanatorily relevant information is whatever information is relevant to a given because-claim on that occasion.

In this paper, I will defend UNIFORMITY. As it stands, UNIFORMITY speaks somewhat vaguely of "explanation". It is vague, because it does not say what a theory of explanation is a theory *of*. There are several different phenomena in the vicinity of our explanatory practices that are candidates. So in order to obtain a thesis precise enough to argue about, we need to conjoin UNIFORMITY with a claim about what explanations are. Two main candidates are important for this paper. On the one hand, there is the practice of asking and answering why-questions. On the other, there is the evaluation of because-claims for truth and falisty. These two are related, but it is a non-trivial matter which one we choose. Let me mark the options explicitly.

[WHY EXPLANATIONS] A theory of explanation is a theory of the interpretation of why-questions (a semantics for why-questions).

[BECAUSE EXPLANATIONS] A theory of explanation is a theory of the interpretation of because-claims (a semantics for because-claims).

If WHY EXPLANATIONS is true, then UNIFORMITY stands and falls with the interpretation of why-questions. So if there is a philosophically illuminating account of their interpretation, UNIFORMITY holds. On the other hand, if BECAUSE EXPLANATIONS is true, UNIFORMITY stands and falls with the interpretation of because-claims.

Van Fraassen endorses WHY EXPLANATIONS, and hence his challenge to UNIFORMITY takes the form of a semantics for why-questions according to which there is no philosophically illuminating account of their interpretation.

Against van Fraassen, I'll argue that WHY EXPLANATIONS is false. Hence, whether there is a philosophically illuminating account of the interpretation of why-questions is irrelevant to whether we should accept UNIFORMITY or not (sections 2-4). My argument against WHY EXPLANATIONS is this: a theory of explanation needs to address certain phenomenon. The example I'll focus on, along with van Fraassen, is the asymmetry of explanation. But a theory of why-questions does not account for the relevant features of that phenomenon, whereas a semantics for because-claims does. Hence, we should reject WHY EXPLANATIONS and accept in its stead BECAUSE EXPLANATIONS. In this paper, I do not take a stand on whether van Fraassen's semantics for why-interrogatives are correct. I only claim that, even if they were correct, they would not threaten UNIFORMITY.

I'll then give two arguments (section 5) for UNIFORMITY. The first is positive: I'll argue that the patterns in our acceptability judgments for because-claims are best accounted for by hypothesizing that they are not context-sensitive in a way that threatens UNIFORMITY (section 5.1). The second argument is negative. I'll argue that the cases van Fraassen presents do not give a proponent of UNIFORMITY reason to relinquish her position (section 5.2). So all of the evidence usually marshalled in this debate in fact points in the same direction.

2 Semantics for Questions

Since I am making a very general claim about semantics for questions, it will pay to be quite explicit about what such semantics look like. It will be easiest to introduce them by considering a simpler example than why-questions, who-questions. They exhibit all of the features we will need, so they serve to illustrate the basic features of a semantics of questions and the issues addressed by such semantics.

2.1 SEMANTICS FOR INTERROGATIVES

Let me start by fixing some terminology. *Indicative sentences* are sentence types. On particular occasions, we token the type and thereby express *propositions*, which are abstract objects. Which proposition is expressed by a sentence-type on an occasion may vary with context. We need not say anything further about what propositions are: they may equally be sets of worlds, Russellian ordered n-tuples, or Fregean thoughts.

We focus on propositions, because they allow us to describe various semantic relations between sentences, such as entailment and incompatibility. A semantic account for indicative sentences is supposed to give us an illuminating account of how the proposition expressed by a sentence is determined as a function of the meanings of its constituent parts, together with their mode of combination, and perhaps together with certain relevant features of the situation in which the sentence is produced.

Interrogative sentences—interrogatives for short—are also sentence types. On particular occasions, we token the type and thereby express *questions*, which are abstract objects.[4] So interrogatives stand to questions as sentences do to propositions.[5]

Unlike propositions, questions are neither true nor false, so we should not identify questions with propositions. Hence, we need a hypothesis about what sort of abstract objects they are, chosen with an eye to being able to define relevant relations between questions, as well as between questions and propositions that

[4]This means, alas, that the title of my paper, as well as most of the text before this note, exhibits conceptual confusion: to speak of the semantics of questions is like speaking of the semantics of propositions. Neither questions nor propositions are the sorts of things one gives semantics for. Now that the terminology has been fixed, I will abide by it.

[5]I think this terminology is more perspicuous than talking about questions (= my interrogatives) and logical questions (= my questions), which can be found in, e.g., Temple (1988) and Van Fraassen (1980).

are their answers. Many theorists identify a question with a set of propositions, either the set of all of its possible answers, or its true answers. I will assume the former option.[6] Semantics for interrogatives seek to give an illuminating account of how the question expressed by an interrogative is determined as a function of the meanings of its constituent parts, together with their mode of combination, and perhaps together with certain relevant features of the situation in which the interrogative is produced.[7] Given the identification of questions with the set of their answers, semantics for interrogatives tell us how the conditions on being an answer to the question are determined by these factors.

Consider an example.

(1) Is it raining?

Semantics for (1) need to predict that the following propositions are all and only the answers to (1).

(2) a. It is raining.

b. It is not raining.

[6]See Groenendijk and Stokhof (1997), Hamblin (1958), Hamblin (1973), and Karttunen (1977). Van Fraassen (1980) follows them in this assumption.

[7]One question one needs to settle is whether in a context, the set of answers is mutually exclusive and jointly exhaustive of the possibilities open in the context. In that case, a question *partitions* logical space, and we can define complete answers as ones that single out exactly one such partition. We can also define less than completely answers in terms of such a partition. This will make a difference when it comes to embedded interrogatives, but the issues at stake do not matter for me here. For a lengthy exposition of the issues, see Groenendijk and Stokhof (1997).

This is not controversial, given the assumptions so far.

2.2 Semantics and Pragmatics for Interrogatives

However, controversy is just around the corner. Consider (3).

(3) Who is Bill Clinton?

Given our assumptions, only propositions of the form (4) qualify as answers to (3).[8]

(4) **Bill Clinton is —.**

However, it is not clear that all such propositions are answers to the question expressed by (3) on a particular occasion of use. Let us say that a *response* is the speech act we perform after a question has been asked. That includes saying things like 'I don't know' or 'I won't tell you.' Uncontroversially, even among the responses that are more responsive, some are better than others. In some situations, saying that Bill Clinton is the forty-second President is a good response, but saying that Bill Clinton is the guy at the bar is not. In other situations, the reverse might be true. Here are two hypotheses about this contrast. Either, both responses express answers, one better than the other; or only one of them does.

[8]Propositional schemata will be bold; sentence schemata will be in corner quotes.

On the former description, the semantics of (3) place fewer constraints on propositions to count as answers, and the difference between the two replies is due to factors beyond the constraints imposed by the question itself—one may be more helpful than the other, for example. On the latter description, the semantics of (3) places more constraints on propositions to count as answers. Settling which of these two descriptions is true is distinguishing the semantics and pragmatics of an interrogative.

What turns on this difference? For one, interrogatives do not just appear as free-standing sentences. They also appear embedded in larger constructions, such as (5), where they make a systematic contribution to the truth-conditions of these larger constructions.

(5) John knows who Bill Clinton is.

Suppose that a knowledge ascription with an embedded interrogative like (5) is true if and only if the subject knows at least one of the answers to the question expressed by the interrogative. Then how we draw the distinction between the semantics and pragmatics of (3) influences what we predict as the truth-conditions of (5). The fewer constraints the semantics of (3) places on its answers, the easier it is for (5) to be true. For example, if we say that all that is required in order for a proposition to be an answer to (3) is that it instantiate (4), then it is sufficient for

John to know who Bill Clinton is that John knows that Bill Clinton is the forty-second President. Indeed, it might be enough for John to know who Bill Clinton is that John knows that Bill Clinton is a person. Alternatively, if the semantics places further constraints on possible answers, then in different situations, different bits of knowledge will be required in order for John to know who Bill Clinton is.[9] Casting our eye back at the sentences exhibiting the asymmetry, some explicitly embed a why-interrogative, such as 'that the shadow is 40ft long explains *why the flagpole is 30ft high.*' Similar issues to the ones I just mentioned arise there.

So quite generally, when we think about the semantics and pragmatics of questions, we try to do the following. We have intuitions about what is a good or bad response to an interrogative on a particular occasion. We then try to separate those responses that actually give answers from those that do not, and explain our intuitions in terms of that distinction. Some responses may be bad, even though they express answers, some responses may be good even though they do not express answers.[10] Van Fraassen presents a theory that addresses this issue for why-questions: which constraints does a proposition have to satisfy to count as an answer to the question expressed by an interrogative on a particular occasion?

[9]Boër and Lycan (1986) make the argument that the semantics of who-questions like (3) places relatively demanding answerhood conditions on propositions by pointing to knowledge ascriptions like (5). They argue that, on different occasions, we have differing intuitions about the truth-value of (5), even when we hold constant the propositions the subject of the ascription knows. That suggests that different propositions count as answers to (3) on different occasions. Needless to say, this view has not gone unopposed.

[10]As an example of the latter, consider this exchange: "Who is Bill Clinton?"—"Think government!" The response may well jog the speaker's memory, and thus be a good response, without expressing an answer to the question.

2.3 Van Fraassen's Theory of Why-Interrogatives

A little bit more terminology to start out. A *why-question* is a question that can be expressed by an interrogative of the form $\ulcorner$Why is it the case that p?$\urcorner$. Its *presupposition* is the proposition that p. We will take this interrogative as canonical. Short forms, such as simply uttering 'Why?' have to be interpreted as short for a canonical interrogative. The object of van Fraassen's theory is to give necessary and sufficient conditions for a proposition to be an answer to a why-interrogative when evaluated with respect to a context C.

Given our stipulations about what a why-question is, and what the presupposition of a why-question is, we can give the following necessary condition:

[PRESUPPOSITION] A proposition that X is an answer to the question expressed by the interrogative $\ulcorner$Why is it the case that $p\urcorner$ in C *only if*

the proposition that X is of the form ***p* because *q***.

PRESUPPOSITION is a fancy way of saying that all answers to our canonical interrogatives have to be propositions of the form ***p* because *q***. For example, all answers to (6a) have to be of the form (6b).

(6) a. Why did the chicken cross the road?

b. **The chicken crossed the road because —.**

You might think at this point that it is odd for van Fraassen to help himself to propositions of the form ***p* because *q***. That is not a problem for van Fraassen: he can give some account of because-claims, and he can give it in terms of answerhood conditions to why-interrogatives. I'll point out below how he can do so non-circularly. To see that, we need to have the remainder of van Fraassen's account before us.

According to van Fraassen, we need more conditions than just PRESUPPOSITION. Sometimes, we reject responses to a why-interrogative that satisfy PRESUPPOSITION, and we restate our question to convey what kind of information we want. These restatements may have the form ⌜Why is it the case that p rather than that r?⌝ We draw a certain contrast. To give van Fraassen's example, we might elaborate on (7) by any of (7a)-(7c).[11]

(7) Why did Adam eat the apple?

a. Why did Adam eat the apple rather than a pear?

b. Why did Adam rather than the snake eat the apple?

c. Why did Adam eat the apple, rather than juggle it?

Being told that Adam ate the apple because he was hungry might be a good reply if we asked (7c), but not if we asked (7a) or (7b).

[11]See Van Fraassen (1980, p. 127).

Van Fraassen suggests that we are tracking semantic facts here. That is to say, the interrogative (7) is semantically ambiguous between at least the three questions expressed by (7a)-(7c). Whether the proposition that Adam ate the apple because he was hungry is an answer depends on which of these questions one expresses. When the contrast is not explicit, as it is not in (7), it is supplied by the context. So we need another necessary condition.

[CONTRAST] A proposition that X is an answer to the question expressed by the interrogative ⌜Why is it the case that p⌝ in C *only if*

the proposition that X is of the form ***p* because *q*** and the proposition that q addresses the contrast salient in C.[12]

I will not have more to say about what it takes for a particular contrast to be salient. Specifying one explicitly in the interrogative makes it salient, as do certain kinds of stress. I also won't have more to say about what contrasts are. The intuitive

[12] I want to emphasize that I am not endorsing van Fraassen's argument. It is an unargued assumption that our intuitions about answers we would reject in one situation but not another are indicative of a difference in the answerhood conditions the question expressed imposes. To see what this assumption amounts to, let me put it in terms of the example involving who-questions: from the observation that sometimes one answer will be accepted to "Who is Bill Clinton?", van Fraassen would directly infer that the interrogative expresses different questions on these two occasions. So van Fraassen would just ignore the other alternative, that the acceptability or unacceptability of the response is due to something other than the answerhood conditions imposed by the question. He ignores the same position in the arguments he gives about the semantics of why-interrogatives. I do not press this point in the text, since I argue that *even if* we grant him everything he says, his account is untenable. For the record, that CONTRAST is a constraint on the semantics of why-interrogatives is more widely held than that RELEVANCE is—see for example Garfinkel (1981), Lipton (1991a), and Lipton (1991b).

But if one wanted to defend a particular view of the semantics of why-interrogatives, one would have to defend the claim that we are tracking semantic facts in our intuitions about (7)—or for that matter, in our intuitions about (8), below. Lewis (1986) questions this step of the argument.

grasp we have from considering examples like (7a)-(7c) should suffice. However, the notion of a proposition's *addressing* a contrast will be important. We can tell that there is a difference between answering "because Adam was hungry" in response to "Why did Adam eat the apple rather than juggle it?" and doing so in response to "Why did Adam eat the apple rather than the pear?" But saying that much does not give us a theory of what it takes for a proposition to address a contrast. I'll return to this issue (section 5.2) once I have presented the rest of van Fraassen's theory and clarified some further issues.

In a conversation, we still reject some propositions as bad responses even though they satisfy both PRESUPPOSITION and CONTRAST. Van Fraassen gives the example of (8), which can be answered either by (8a) or (8b).[13]

(8) Why did the conductor warp rather than retain its shape?

   a. The conductor warped rather than retain its shape because a very strong current went through it.

   b. The conductor warped rather than retain its shape because it was in the Earth's magnetic field.

The contrast between these two sentences is supposed to arise in a situation where it is true both that a very strong current went through the conductor and that the

[13] See Van Fraassen (1980, pp. 141-2).

conductor was in the Earth's magnetic field. These facts remain constant. The intuition is that, depending on what we are interested in, one or the other of (8a) and (8b) is an answer to (8). According to van Fraassen, this, too, tracks a difference in the question expressed by (8) on different occasions, and we therefore need another condition.

[RELEVANCE] A proposition that X is an answer to the question expressed by the interrogative ⌜Why is it the case that p⌝ in C *only if*

the proposition that X is of the form ***p* because *q*** and the proposition that q bears the relevance relation salient in C to the proposition that p and the contrast salient in C.

The idea behind RELEVANCE is that the presence of the Earth's magnetic field is a background condition of the conductor's warping, whereas the surge of current is an efficient cause. The propositions describing these two bear different relevance relations to the presupposition and contrast of (8). Let me introduce a name for the resulting position.

[SEMANTIC APPROPRIATENESS] The proposition that ***p* because *q*** is a *semantically appropriate* answer to the question expressed by a why-interrogative ⌜Why is it the case that p?⌝ in C iff

(i) the proposition that ***p* because *q*** satisfies PRESUPPOSITION; and

(ii) the proposition that q addresses the contrast salient in C and bears the relevance relation salient in C to the proposition that p and the contrast salient in C.

SEMANTIC APPROPRIATENESS summarizes the semantic factors that determine whether a given response expresses an answer. If a semantically appropriate answer is also true, all the semantic hurdles have been cleared for the response to be appropriate *simpliciter*. If a response is inappropriate, even though it expresses a true, semantically appropriate answer, that inappropriateness is due to some other feature either of the answer or of the sentence uttered, such as rudeness, triviality, or not being presented in a language spoken in the context. In the usual case, a response will be appropriate *simpliciter* only if it expresses a semantically appropriate answer.[14]

Many commentators complain that van Fraassen is silent on any further characterization of the relevance relation.[15] But this is not an oversight. He is silent on a further characterization, because he does not think that we *can* characterize it further. Potentially, anything could be a relevance relation, and so the speaker interests that fix which relation is the relevance relation on a particular occasion of use are not constrained.

[14]A case where this fails: "Why do the tides exhibit the pattern they do?"—"Think Newton!" If the response prods the memory of the person asking the question, the response may have been appropriate, even though it does not express a semantically appropriate answer. Consider also 'I don't know' along with the example in note 10.

[15]See for example: Grimes (1987, p. 89), Hitchcock (1996, p. 399), Kitcher and Salmon (1987, p. 318), Sandborg (1998).

To see how this works exactly, let me consider van Fraassen's most important case study, the asymmetry of explanation.

3 Asymmetry of Explanation

The asymmetry of explanation is best introduced by example, since the debate is about how to characterize it abstractly. Here is a famous one.[16]

[FLAGPOLE] A flagpole is standing in a field. As the sun moves through the sky, the flagpole casts a shadow of varying length. Right now, the sun is shining, and the pole casts a shadow. The pole is 30ft high and the shadow is 40ft long.

Observe the following contrasts.

(9) a. Why is the shadow 40ft long? — Because the flagpole is 30ft high.

b. Why is the flagpole 30ft high? — Because the shadow is 40ft long.

(10) a. That the flagpole is 30ft high explains why the shadow is 40ft long.

b. That the shadow is 40ft long explains why the flagpole is 30ft high.

(11) a. The shadow is 40ft long because the flagpole is 30ft high.

[16] Hempel discusses it in Hempel (1965a), where he credits it to Bromberger. After that classic discussion, every account of explanation has had to come to terms with it.

b. The flagpole is 30ft high because the shadow is 40ft long.

In each pair, the first member is acceptable, the second not. The problem of the asymmetry of explanation in the first instance is to say what makes for the difference between the members of each pair. In the second instance, it is to deal with other explanatory asymmetries, which may or may not be susceptible to a similar treatment.

The contrasts between the members of the pair are called "asymmetries" for historical reasons. One of the central questions twentieth-century empiricists wanted to address was how much metaphysics our best scientific theories and our best scientific practice commit us to. Empiricists wanted to give a very minimal answer to that question. On their view, we are committed to observable objects and properties, and to relations between them that make for well-justified predictions. An example of such a relation might be the relation of statistical relevance.[17]

To show that the minimal answer is correct, empiricists tried to capture all the relevant features of scientific practice—including the practice of giving explanations—using only the tools they limited themselves to. A crucial part of this empiricist project is the DN-account of explanation I mentioned earlier. On this view, we have an explanation of the fact that p just in case we can deduce a statement of

[17]Indeed, this is the relation most empiricists focused on, perhaps with the added requirement that the statistical correlation had to be such as to make for justified predictions. In the terminology of Goodman (1983), the correlation had to be projectible.

that fact from law-statements and statements about some initial conditions. But this account faces a systematic problem. A derivation of the height of the flagpole from statements about the length of the shadow and the angle of the sun satisfies its criteria for being an explanation just as much as a derivation of the length of the shadow from statements about the height of the flagpole. The problem is that the laws used in the derivation—trigonometry—are symmetrical, but explanation is not. Only one of the corresponding explanations is acceptable, the one that explains the length of the shadow by appeal to the height of the flagpole. At times, Hempel seems to have been inclined towards dismissing the asymmetry as not showing us anything of importance about science, and hence not requiring a mature philosophy of science to account for it.[18] But that is no longer viable.[19]

In response, many philosophers have taken the explanatory asymmetry to reflect an asymmetry in the phenomena, specifically an asymmetric causal relation between the flagpole and the shadow it casts.[20] If they are right, our best scientific theory and practice commits us to more metaphysically more loaded relations than empiricists wanted to countenance. Against this majority view, van Fraassen wants to give an alternative account of the asymmetry, according to which they

[18]Hempel seems to have favored this account at one time. See his discussion of the pendulum at Hempel (1965a, pp. 352-3).

[19]Salmon has explicitly engaged this attitude in his (1998), where he argues against it.

[20]This conclusion is drawn by all of the theorists who say that all explanation is causal explanation, such as Hausman (1998), Lewis (1986), Salmon (1984), Strevens (2004), and Woodward (2003). Alternatively, one might hold that the metaphysical asymmetry in the subject situation consists in the presence of asymmetric laws. This is the approach taken in Bromberger (1992c).

do not force us to recognize metaphysically more loaded relations than statistical correlation.[21]

3.1 VAN FRAASSEN ON THE ASYMMETRY OF EXPLANATION

The key move van Fraassen makes is to say that the bad members exhibiting the asymmetry are bad because the salient relevance relation is not satisfied, and that this relevance relation is not metaphysically loaded. Let me explain this in more detail.

Recall the pairs of sentences exhibiting the asymmetry of explanation in FLAGPOLE.

(9) a. Why is the shadow 40ft long? — Because the flagpole is 30ft high.

b. Why is the flagpole 30ft high? — Because the shadow is 40ft long.

(10) a. That the flagpole is 30ft high explains why the shadow is 40ft long.

b. That the shadow is 40ft long explains why the flagpole is 30ft high.

(11) a. The shadow is 40ft long because the flagpole is 30ft high.

b. The flagpole is 30ft high because the shadow is 40ft long.

[21] As he says in introducing that theory, "I shall offer a model of this aspect of scientific activity [i.e., explaining] in terms of why-questions, their presuppositions, and their context dependence. This will account for the puzzling features (especially asymmetries and rejection) that have been found in the phenomenon of explanation, while remaining compatible with empiricism."(Van Fraassen, 1980, p. 97).

Van Fraassen first applies the account of why-questions to the exchanges in (9), and then proposes interpretations for the sentences in the remaining pairs to show that they inherit the asymmetry of (9). The main idea is that either explicitly or implicitly, the remaining sentences contain embedded why-interrogatives.

Since (9a) and (9b) involve different questions, the questions will impose different answerhood conditions. Moreover, since the response in (9a) is appropriate, it must express a semantically appropriate answer. But suppose now that (9b) is bad, because the proposition expressed by the response does not satisfy the answerhood conditions imposed by the question asked. So *that* proposition is not a semantically appropriate answer to the question. The difference in semantic appropriateness explains the felt difference in appropriateness *simpliciter*. That is the first part of the account.

Clearly, if the answer in (9b) is semantically inappropriate, that must be due to its failing to satisfy RELEVANCE. Even if we made the answer so that it explicitly satisfied PRESUPPOSITION and CONTRAST, it would still be bad. "Why is the flagpole 30ft high rather than some other height?—The flagpole is 30ft high rather than some other height because the shadow is 40ft long" is just as bad as (9b).[22] But we might think at this point that RELEVANCE is not satisfied because the length of the shadow is not *causally relevant* to the height of the flagpole. If that was all there was to van Fraassen's proposal, it would hardly merit our attention. It

[22]That (9b) is bad because it fails to satisfy RELEVANCE is also van Fraassen's diagnosis. See Van Fraassen (1980, p. 130).

certainly would not bring any anti-metaphysical conclusions with it. I will address this worry at length below. For now, I want to finish out van Fraassen's treatment.

Here is how we can extend the account to the remaining pairs. A plausible set of truth-conditions for (10a) and (10b) might be these.

- A sentence of the form ⌜that q explains why it is the case that p⌝ is true in a context C if and only if the proposition that p because q is a true, semantically appropriate answer to the question expressed by ⌜why is it the case that p?⌝ in C.[23]

In that case, the asymmetry between (9a) and (9b) carries over directly to (10a) and (10b). Since the proposition that the shadow is 40ft long because the flagpole is 30ft high is a semantically appropriate answer to the question expressed by the embedded interrogative, (10a) is true. Likewise, since the answer in (9b) is not semantically appropriate, (10b) is false.

Van Fraassen's treatment extends to the two remaining because-claims, as well. He just has to assume that the unacceptability of the because-claims is due to the fact that some salient answerhood conditions are not satisfied. In that case, the treatment goes like this.

[23]I am thus disagreeing with Grimes, who claims that the schema " 'Its being the case that E is a scientific explanation for why it is the case that P in contrast to X' [is one] where [...] there is no reference to any question or answer."(Grimes (1987, p. 88)) Grimes wants to conclude that the theory of why-questions is therefore irrelevant even to sentences like (10a) and (10b). But an embedded question is a question, and hence the schema does make reference to a question. Grimes may have something else in mind; perhaps he wants to deny that the semantics of interrogatives on their own carry over to sentences in which they are embedded. But he does not argue for any such thesis.

(11) a. The shadow is 40ft long because the flagpole is 30ft high.

b. The flagpole is 30ft high because the shadow is 40ft long.

In order to deliver the prediction that (11a) is true and (11b) false, we just need one more assumption. The assumption is that, when we evaluate these two sentences, the salient contrast and relevance relation are the same as when we evaluate (9a) and (9b), respectively. Then the account runs thus: by assumption, that the flagpole is 30ft high satisfies the contrast and relevance salient when we ask why the shadow is 40ft long. Since these conditions are also the salient ones in the context in which we evaluate (11a), (11a) is true. Again, by assumption, that the shadow is 40ft long does not satisfy the answerhood conditions imposed by the question why the flagpole is 30ft high. Since *those* conditions are salient in the context in which we evaluate (11b), (11b) is false. That explains the asymmetry.

3.2 What is the Relevance Relation?

As I pointed out in presenting van Fraassen's response to the asymmetry, for all we have said so far, the relevance relation that isn't satisfied by the bad members of each pair is just a causal relevance relation. That is to say, the answers are bad because they do not cite causally relevant factors. At times, van Fraassen seems to endorse this diagnosis of the case, as well. Consider his famous discussion of the tower and the shadow.[24] That story is supposed to provide a context in which

[24]See Van Fraassen (1980, pp. 132-4).

the bad because-claim is good. When the tower was built, the designer wanted to ensure that the shadow would cover a certain spot on the ground at sunset. For that, the shadow had to have a certain length, and that in turn required the tower to be a certain height. Given this background information, we are supposed to find it reasonable to say (12).

(12) The tower is 100ft high because the shadow is 200ft long.

I think that (12) still sounds odd, even in the situation van Fraassen describes. But I want to draw attention to a different point.

Suppose that (12) was acceptable, and that it was acceptable because a different relevance relation is salient in this context than is salient in the usual contexts. It may still be the case that all the relevance relations involved—the one salient in the usual contexts and the one salient in the context van Fraassen describes in the story—are metaphysically loaded causal relations. Indeed, that is van Fraassen's own gloss on the example when he introduces it.

> [In this story,] the relevance relation changes from one sort of efficient cause to another, the second being a person's desires.[25]

That is, depending on the salient relevance-relation the sentence "the flagpole is 30ft high because the shadow is 40ft long" expresses a falsehood or a truth. But

[25] Van Fraassen (1980, p. 132).

either of the salient relevance relations is metaphysically loaded, since they are both causal relations. If that were right, van Fraassen would have just refuted his own project.[26]

So we should reject this gloss on the example. Instead, relevance relations need to be drawn from a different set of relations, one not including metaphysically loaded causal relations, if we are to preserve the metaphysical innocence van Fraassen aims for. Let's focus on the because-claims exhibiting the asymmetry.

(11) a. The shadow is 40ft long because the flagpole is 30ft high.

b. The flagpole is 30ft high because the shadow is 40ft long.

Let's call them the good and the bad because-claim, respectively. All van Fraassen needs in order to account for the asymmetry is that, in the contexts in which we ordinarily evaluate the two claims, there is some asymmetric relation that holds

[26]Compare in this context also van Fraassen's earlier summary of his position:

> In the case of causal explanation, the *explanation* consists in drawing attention to certain ('special', 'important') features of the causal net. (Van Fraassen, 1980, pp. 124f)

And again later in discussing the relevance-relation.

> For example, suppose you ask why I got up at seven o'clock this morning, and I say 'because I was woken up by the clatter the milkman made'. In that case I have interpreted your question as asking for a sort of reason and my word 'because' indicates that the milkman's clatter was that sort of reason, that is, one of the events in what Salmon would call the causal process. (Van Fraassen, 1980, p. 143)

I'm not sure how to interpret these remarks. I will assume that metaphysical innocence is the overriding aim of the account, so I will disregard these comments. If I am thereby making a historical mistake, I should be taken to be discussing the philosopher van Schmaassen, for whom metaphysical innocence really is the overriding consideration.

between the shadow and the flagpole, but not vice versa, and that the acceptability of the because-claims depends on the obtaining of that relation. There are many potential candidates: being shorter than, to name just one. If that is the salient relevance relation, then we predict that the good because-claim is acceptable, and the bad one not, since the length of the flagpole satisfies the relevance relation with respect to the shadow, but not vice versa.

This particular proposal faces an obvious objection: there are many true because-claims where *that* relevance relation is not satisfied. Consider an ordinary pendulum. It's true that the pendulum's period is 2 seconds because the length of its string is 1 meter. But the length of the string and the length of the period don't even have the same dimension, and hence cannot satisfy a relation of being shorter than.

The objection cuts any ice at all *only if* the same relevance relation has to be satisfied by both because-claims, the flagpole and the pendulum. But that assumption fails if the relevance relation is contextually determined. In that case, it may be true that the good because-claim about the flagpole is true because it satisfies one relevance relation, and the because-claim about the pendulum is true because it satisfies another. Van Fraassen's idea, thus, is to say that the context-dependence of because-claims is so pronounced that we cannot rule out a wide diversity of relevance relations that are satisfied in different cases, all of which are metaphysically innocent.[27]

3.3 Why Explanations and Because Explanations

If this last claim is correct, then there is no philosophically illuminating semantic theory for why-interrogatives. Whether a proposition is an answer to the question expressed by a given interrogative (on a particular occasion of use) depends on unconstrained features of the context.[28]

Suppose for now that van Fraassen is right about the prospects for a semantics of why-interrogatives. Does that mean that UNIFORMITY fails? Not immediately. As I said in the introduction, the crucial question for van Fraassen is whether a theory of explanation is a theory of why-questions. He certainly endorses WHY EXPLANATIONS:

> An explanation is not the same as a proposition, or an argument, or

[27]That is the best reading available, I think. As the passages I cite in note 26 show, van Fraassen at times seems completely comfortable countenancing causation. But on the reading I propose, we retain metaphysical innocence while also according context-dependence a prominent place in the argument.

[28]Let me briefly point out that this does not show that, on a particular occasion of use, very specific information is required in order for a because-claim to be true. And for all that van Fraassen has said, there may be patterns of relevance relations that are salient in a whole host of contexts.

Kitcher and Salmon (1987) seem to me to overlook this point in their influential criticism of van Fraassen's account. They point out that an astrological relevance relation is not one that will ever be salient for many of us. They infer from this observation that such an astrological relevance relation is ruled out on the grounds that it is not an *explanatory* relation. But that is precisely what is at issue between Kitcher and Salmon and Van Fraassen. Both agree that we do not find ourselves in contexts in which astrological relevance relations are salient. Kitcher and Salmon claim that that is the case because the concept of an explanation constrains relevance relations. Van Fraassen claims that the relevance relation is constrained by something other than the interest in explanation.

list of propositions; it is an *answer.* An explanation is an answer to a why-question.[29]

But I not want to argue that WHY EXPLANATIONS is in fact false. It fails, because the issues that made us care about a theory of explanation in the first place are not properly reconstructed as issues about the answerhood conditions of why-interrogatives. I will argue this point for the special case of the asymmetry of explanation.

Answerhood conditions are different from truth-conditions. The example of who-questions and their answers makes this point uncontroversially. Consider again (13), used on a particular occasion, and three things one might say in return, (13a), (13b), and (13c).

(13) Who is Bill Clinton?

a. Bill Clinton is the Forty-Second President.

b. Bill Clinton is the guy at the bar.

c. Bill Clinton is my brother.

It may well be the case that on an occasion, only the proposition expressed by (13a) is an answer to the question expressed. Nonetheless, both (13a) and (13b)

[29] Van Fraassen (1980, p. 134).

are obviously true. On the other hand, (13c) is simply false. So while both (13b) and (13c) are less than fully acceptable in the context, their unacceptability has different sources. A proposition's unacceptability that is due to failing to satisfy some salient answerhood conditions is compatible with the proposition's truth.

With that in mind, return to the bad because-claim in the case of the flagpole, (11b).

(11b) The flagpole is 30ft high because the shadow is 40ft long.

The question for us is whether we should think that its unacceptability is due its being false or its failing to satisfy salient answerhood conditions. The latter option is compatible with both because-claims' being true. So far, I've only spoken of the unacceptability of the bad because-claim in order to remain neutral on this issue.

Now I want to argue now that the bad because-claim is bad because it is *false*. I urge this *not* on the grounds that this is intuitively obvious. The notion of truth employed here is a semi-technical notion, and our intuitions do not count for much on this matter.[30] However, there are some more theoretically motivated arguments for accepting this position.

Two main competitors need to be ruled out. The first is van Fraassen's preferred account, based on answerhood conditions. The bad because-claim is unacceptable because there are answerhood conditions salient when we evaluate it that

[30]That is because the notion of truth is connected to other notions like compositional meaning, presupposition, and assertion.

the claim fails to satisfy. Perhaps we always evaluate because-claims as answers to an implicit why-question. The other is a broadly Gricean account, according to which uttering the bad because-claim is bad because it violates some principles governing cooperative conversation.[31]

First, the asymmetry of explanation arises in thought just as much as it does in language. Not only do we not want to *say* that the flagpole is 30ft high because the shadow is 40ft long, we also do not *believe* it. This observation disqualifies the Gricean strategy, because Gricean principles govern what one can say. They are inapplicable to thought, and hence cannot be used to give an account of why we are not inclined to believe the bad members of each pair. Perhaps answerhood conditions can be salient in thought as much as in language, so this objection may not tell completely against van Fraassen's position. However, the next one seems to me to be fatal to both.

If a sentence is true but unacceptable, its denial is not acceptable either. This is a straightforward prediction. If a sentence is true, its denial is false. Falsehoods are usually unacceptable, and even if they are acceptable, they are degraded. So if the denial of the bad members of each pair is acceptable, that gives us good reason to think that the original was false. To that end, observe the acceptability of (14).

[31] As they are presented in Grice (1991). The distinction between van Fraassen's preferred account and the Gricean one is not sharp. One of the principles governing conversation requires speakers to make useful contributions. When a question is asked and the response does not express an answer, the response fails to make a useful contribution to the conversation and is therefore less than fully acceptable. I distinguish van Fraassen's and Grice's account in the text more for ease of exposition, and we'll see that they fail for the same reasons.

(14) It is not the case that the flagpole is 30ft high because the shadow is 40ft long.

These two arguments suggest that any account of the asymmetry has to predict that the bad members of the pairs exhibiting the asymmetry are false, not just true propositions that do not satisfy the answerhood conditions imposed by some why-question, or true propositions whose assertion fails to conform to Gricean principles of conversational conduct.

3.4 Challenging Uniformity on Truth

If I am right in arguing that the bad because-claim is bad because it is false, then van Fraassen's argument against UNIFORMITY falls flat. It fails, because WHY EXPLANATIONS is false, and thus because the semantics of why-interrogatives are irrelevant to UNIFORMITY.

A theory of explanation is supposed to cover phenomena like the asymmetry of explanation, and the arguments I have just gone through show that doing so requires giving a semantic theory for because-claims, not a semantic theory for interrogatives. So at least in the first instance, we should accept BECAUSE EXPLANATIONS. Joined to UNIFORMITY, the proposition under discussion is therefore that because-claims have a uniform semantics, and that we can give a philosophically illuminating account of these semantics. Without further ado, van Fraassen's conclusion does not conflict with *that* claim, since his conclusion concerns the semantics of why-interrogatives.

The qualification 'without further ado' is important. We can, of course, connect van Fraassen's theory of why-interrogatives to a semantic theory for because-claims. The key idea is to say that the truth of a because-claim in a context depends on whether it satisfies salient answerhood conditions. I think van Fraassen in fact endorses this thought. More explicitly, the truth-conditions might go like this.[32]

[TRUTH-CONDITIONS FOR 'BECAUSE']

$\ulcorner p$ because $q\urcorner$ is true in a context C iff

(i) it is true that p; and

(ii) it is true that q; and

[32]Van Fraassen's statement of the truth-conditions are spread throughout Van Fraassen (1980), so let me give some textual support for this interpretation. I take it that the first two conditions, the truth of the propositions that p and that q, are uncontroversial.

The requirement for statistical relevance is given here: "explanatory factors are to be chosen from a range of factors which are (or which the scientific theory lists as) objectively relevant in certain special ways—but that the choice is then determined by other factors that vary with the context of the explanation request. To sum up: no factor is explanatorily relevant unless it is scientifically relevant; and among the scientifically relevant factors, context determines explanatorily relevant ones."(Van Fraassen (1980, p. 126) The contextual factors are the ones determining semantic appropriateness. Van Fraassen elaborates on what he calls the 'scientifically relevant' factors here: "*The context*, in other words, *determines relevance* in a way that goes well beyond the statistical relevance about which our scientific theories give information."(Van Fraassen (1980, p. 129) This shows that to be scientifically relevant is to be statistically relevant, and that justifies my attribution of the third requirement of statistical relevance.

Finally, van Fraassen gives the fourth condition here: "in my opinion, the word 'because' here [in a sentence of the form $\ulcorner$because $A\urcorner$] signifies only that A is relevant, in this context, to [a contextually given] question." (Van Fraassen, 1980, p. 143)

(iii) the proposition that q addresses the contrast salient in C and bears the relevance relation salient in C to the proposition that p and the contrast salient in C.

The reference to the context C should be understood as part of the *character* of 'because', rather than its *content*. To illustrate the difference, the character of 'I' picks out the speaker of the context. On an occasion of use, the content of 'I' is just the object. The reason to draw this distinction is to accommodate two insights. On the one hand, there is clearly a rule that determines what 'I' picks out on an occasion of use—the speaker. This is captured in the expression's character. On the other hand, the two expressions 'I' and 'the speaker' embed very differently in complex sentences. For example, it is false that if the speaker was Bill Clinton, then I would be Bill Clinton, even though it is true that if the speaker was Bill Clinton, then the speaker would be Bill Clinton. This is captured by the expression's content.[33]

The reference to salient contrasts and relevance relations needs to be part of the character to capture the embedding facts. It seems obviously false to say: Adam ate the apple because he was hungry, but if a different contrast was salient, Adam wouldn't have eaten the apple because he was hungry. The truth-conditions as stated right now avoid predicting that this is true.

These truth-conditions also are not objectionably circular, since CONTRAST

[33]For an extended discussion of this, see Kaplan (1989). Arguments concerning the embedding behavior of context-sensitive expressions can be found at Kaplan (1989, pp. 498ff).

and RELEVANCE impose conditions on being an answer that are not themselves put in terms of the truth of because-claims. And they make because-claims context sensitive, since their interpretation depends in part on the context with respect to which they are evaluated.

If these truth-conditions for 'because' are correct, the unconstrained context-sensitivity of why-interrogatives carries over directly to the interpretation of because-claims. And in that case, van Fraassen's theory really does conflict with UNIFORMITY, even if this latter claim is interpreted as I have suggested—as about the semantics of because-claims. So the thesis I want to consider now is whether van Fraassen's semantics for 'because' are acceptable.

4 The Irrelevance of Why-Questions

The first point to make is that, once we focus on the semantics of because-claims, the semantics of why-interrogatives are irrelevant. What matters is whether because-claims are context-sensitive in a certain way, and we cannot infer that they are from the premise that the answerhood conditions of why-interrogatives are context-dependent in this way.

Roughly, the point is this. Van Fraassen predicts that the bad because-claim (11b) is false by saying that it fails to satisfy a contextually salient relevance condition. He makes the further claim that the truth of the because-claim is sensitive to this relevance condition because it is an answerhood condition imposed by a salient why-question, and the truth of the because-claim is sensitive to whether it satisfies the answerhood conditions of a salient why-question. But this theory is equivalent to a theory that writes sensitivity to the salient relevance relation directly into the truth-conditions of because-claims, without the detour via a salient question.

Let me put this point in a slightly different form. Consider the following two theories.

[T_1] (i) *p because q* is true in C iff

- it is true that p, and
- it is true that q, and
- the propositions that p and that q stand in relation R.

(ii) ⌜Why is it the case that p?⌝ is answered by a proposition X in C iff

- X is of the form ***p* because *q***

[T_2] (i) ⌜p because q⌝ is true in C iff

- it is true that p, and
- it is true that q, and
- the proposition that q answers the question why it is the case that p that is salient in C.

(ii) ⌜Why is it the case that p?⌝ is answered by a proposition X in C iff

- X is of the form ***p* because *q***; and
- The propositions that p and that q stand in relation R.

T_1 and T_2 make exactly the same predictions about the truth-conditions of because-claims, and the true, semantically appropriate answers to why-interrogatives.

That means that what matters for our purposes is whether because-claims are context-sensitive in the way that van Fraassen suggests. Whether that context-sensitivity is mediated by why-questions or not is irrelevant. The crucial feature of the case that makes why-questions irrelevant to the asymmetry of explanation is the fact that we are accounting for the *falsehood* of a because-claim. If we wanted to account for the unacceptability of a claim that is nonetheless true, looking to why-questions is at least a very promising option worth exploring. But anytime we want to account for the falsehood of some because-claim, why-questions are irrelevant.[34]

[34]Perhaps one qualification of this claim is in order. Depending on the details, the two theories T_1 and T_2 I consider in the text might make different predictions about the syntax of because-claims. Here is a view of how claims about context-dependence and claims about syntax interact. On this view, a sentence expresses a proposition by being interpreted by a compositional interpretation function. Given the same input sentence, the function returns the same output proposition. The representation of a sentence as input to this function is called the sentence's logical form or LF. For example, sentences that exhibit quantifier scope ambiguity, such as "every boy likes a teacher" is associated with two distinct LFs. And if a sentence expresses different propositions when evaluated with respect to different contexts, that requires that the sentence's LF has an element that is assigned a denotation by the context.

On such a view, the two proposals might have different implications for the LFs of because-claims. According to T_1, there is an element that needs to be assigned some relation R as denotation. According to T_2, there is an element that needs to be assigned a why-question as denotation. One might hope to distinguish the two theories by which of these syntactic theories of because-

5 Arguing for Uniformity

So we have to ask whether because-claims are sensitive to a relevance relation. In the next two sections, I'll offer arguments against this claim. In this section, I'll give positive reason to believe that because-claims are not context-sensitive in the way van Fraassen suggests.

Let's begin by setting aside one kind of context-dependence that is irrelevant to our present concerns. Many because-claims are context-dependent for uninteresting reasons, for example, because they contain pronouns or tensed verbs. This kind of context-dependence is irrelevant for us. We need to focus on forms of context-dependence that are *distinctive* to because-claims.

Van Fraassen suggests that there are two such distinctive kinds, dependence on a salient contrast-class and dependence on a salient relevance relation. The threat to UNIFORMITY arises from the second kind of dependence. All parties to the debate agree that there is a uniform account of addressing a contrast, though theorists may differ on how to spell it out. However, if it is true that potentially anything could be a relevance relation, UNIFORMITY fails. So the crucial claim to be discussed is this.

[VARIABLE RELEVANCE] The truth of because-claims is sensitive to a contextually variable relevance relation.

claims is correct. For arguments about context-dependence that rely on syntactic implications of context-dependence like the ones I have outlined here, see Stanley (2000) and (2002).

Let me comment briefly on VARIABLE RELEVANCE. As I argued earlier, the debate has to be about the *truth* of because-claims, rather than their acceptability that goes beyond truth or falsity. The important feature of the claim is that the relevance relation is *variable*. Everybody can agree that because-claims are sensitive to a relevance relation. Causal theorists of explanation, for example, will suggest that such a relevance relation is causation. Where van Fraassen joins issues with proponents of UNIFORMITY is on the question of how variable this relevance relation can be. Proponents of UNIFORMITY deny, while van Fraassen asserts, that the relevance relation can change from context to context.

The arguments I present in this section are both addressed to VARIABLE RELEVANCE. I'll argue first that there is positive reason to think that there is no variation between contexts in the relevance relation a because-claim is sensitive to.

5.1 THE POSITIVE ARGUMENT

Let us call an *instance of the asymmetry* any pair of because-claims that fit the pattern of (11a) and (11b): one member is of the form $\ulcorner p$ because $q\urcorner$, the other $\ulcorner q$ because $p\urcorner$, and the propositions are statistically relevant to each other. It is a non-negotiable aspect of van Fraassen's position that for every instance of the asymmetry, there is a relevance relation that distinguishes between the acceptable and the unacceptable member of the pair making up the asymmetry. We can marshal evidence to support that in fact, that has to be the *same* relation. That is, we can argue that there is a relevance relation such that for every instance of the asymmetry, *that* relation distinguishes between the members of the pair.

Consider two instances of the asymmetry, such as (11) and (15).

(11) a. The shadow is 40ft long because the flagpole is 30ft high.

b. The flagpole is 30ft high because the shadow is 40ft long.

(15) a. The period of the pendulum is 2sec because the string is 1m long.

b. The string of the pendulum is 1m long because the period is 2sec.

Van Fraassen is committed to the claim that (11a) is true in virtue of satisfying one relevance relation R, while (15a) is true in virtue of satisfying a relevance relation R', where R may or may not be R'. But we can argue that in fact, both (11a) and (15a) satisfy the same relevance relation.

Consider the conjunction of the two because claims (16).

(16) The shadow is 40ft long because the flagpole is 30ft high, and the period of the pendulum is 2sec because the length of the string is 1m.

The conjunction is clearly acceptable. There are two possible reasons for this acceptability. That, by itself, does not show that both conjuncts are true in virtue of satisfying the same relevance relation. It may also be the case that the salient relevance relation changes in the course of evaluating the conjuncts one after the other. That is, we may be dealing with an instance of *accommodation*.[35] Accommodation occurs whenever an assertion is made that is true only if the context satisfies certain conditions, and the context changes so as to make the assertion come out true. Standard examples concern presupposition and the presence of certain standards. For example, one might be in a context in which we accept relatively strict standards for something to have a certain shape, such as a grade school class. In that case, it will be true to say that dice are square, but not that France is hexagonal. However, actually uttering "France is hexagonal" changes the context to one in which the standards are sufficiently relaxed to make the utterance come out true.

The same thing might occur when we consider the conjunction (16). The first conjunct is interpreted with respect to one relevance relation R, the second with respect to another R'. However, supposing that accommodation takes place also makes a prediction. In general, a participant to a conversation can resist the accommodation, usually by denying the claim that would change the standards. In the grade school classroom, the following discourse is a perfectly reasonable way to reject the accommodation.

[35]The notion was first introduced by Lewis (1983).

(17) A: What's square?

B: Dice are square.

A: And what's hexagonal?

B: France is hexagonal.

A: No, it isn't. But this tile is.

This form of resistance is not available for the conjunction (16).

(18) A: The shadow is 40ft long because the flagpole is 30ft high, and the period of the pendulum is 2sec because the length of the string is 1m.

B: # No, it's not the case that the period of the pendulum is 2sec because the length of the string is 1m.

That shows that, at least for the two instance of the asymmetry considered here, a single relevance has to work for both.

I want to draw attention to the form of argument I have developed here. We want to check whether two because-claims are true in virtue of satisfying the same relevance relation or in virtue of satisfying different ones. To determine that, we consider a two-step process. First, we see whether the conjunction of the two claims is acceptable. If it is, we determine whether the context shifts in the course of evaluating the conjunction. If it does not, then we have reason to believe that the because-claims are true in virtue of satisfying a single relevance relation. While it is true that I have chosen the because-claims from instances of the asymmetry, the argument does not actually depend on this fact. All we need are two true because-claims.

In implementing this argument for the particular claims making up (16), I have considered the most widely accepted account of context-shifting: Lewis' account of accommodation. Perhaps the proponent of VARIABLE RELEVANCE—and thus the opponent of UNIFORMITY—will want to supplement her position with a new account of context-shifting that does not make the empirical prediction I have focused on. That is the prediction that one can resist the shift in context by rejecting the assertion that, if unchallenged, induces that shift. But here the burden of proof is clearly on my opponent.

The more because-claims we show to satisfy the same relevance relation by this argumentative strategy, the stronger the case for UNIFORMITY. As an empirical hypothesis, I'll assert that the argument will work for any pair of intuitively acceptable because-claims, and thus, that the prospects for UNIFORMITY are bright.

5.2 THE NEGATIVE ARGUMENT

In the course of presenting his semantics for why-interrogatives, van Fraassen gives various motivating examples. I want to end by considering whether these examples give reason to accept that that because-claims *are* sensitive to a variable relevance relation. I do that, because I want to show that in the debate about UNIFORMITY, all the evidence points in the same direction. We are not left in the uncomfortable position of having some evidence for a claim and some against it.

Specifically, I want to argue that if we hold a semantic theory for because-claims opposed to van Fraassen's, the examples he adduces do not force us to give up that view and move closer to van Fraassen's own position. So I will present this as an exercise in theory choice. Since van Fraassen claims that because-claims are sensitive to two features of the context, contrast-class and relevance relation, theories will have to take a stand on both.

It seems plausible that because-claims exhibit the first kind of dependence van Fraassen mentions, dependence on a salient contrast. Consider the difference between (19a), (19b), and (19c).

(19) a. Adam ate the apple rather than juggle it because he was hungry.

b. Adam ate the apple rather than the pear because he was hungry.

c. Adam ate the apple because he was hungry.

I take it that (19a) is true, (19b) false. Now consider (19c). Depending on which contrast is salient, we might hear (19c) as true or false. That suggests that relative to a context in which the contrast with juggling the apple is salient, (19c) expresses the same proposition as (19a). And relative to a context in which the contrast with eating the pear is salient, (19c) expresses the same proposition as (19b). If that is right, because-claims like (19c) that do not explicitly specify a contrast are context-dependent, since they express different propositions when evaluated with respect to different contexts. That suggests that all theories we choose between should incorporate context-dependence on a salient contrast.

But it is not enough to say simply that, depending on a salient contrast, because-claims without explicitly specified contrast classes express different propositions. We also need to say something about how to interpret this dependence on a contrast-class. Let us say that given a because claim *p rather than X because q*, what matters is for the proposition that *q* to *address* the contrast. In these terms, (19b) is false because the proposition that Adam was hungry does not address the contrast between eating the apple and eating the pear. That much will be common ground between all theories we might consider. But theories will differ on what it takes for a proposition to address a contrast.

All theories about because-claims thus take a stand on two issues. First, they incorporate some account of what it is to address a contrast. Second, they take a stand on whether because-claims exhibit any dependence on a salient relevance relation, and if so, how that relation is constrained. Given the dialectical situation we are in, we should consider van Fraassen's theory along with one other, the one that is most opposed to his. A natural suggestion here is to consider a theory that includes an appeal to causation, and that denies that because-claims exhibit the second distinctive kind of context-dependence, dependence on a salient relevance relation. If we can marshal evidence for or against this competitor, we will have good reason to accept or reject van Fraassen's theory.

The first theory says that a proposition addresses a contrast only if what is described by that proposition is causally relevant to that contrast.[36] On this theory, we account for the asymmetry of explanation by appeal to the contrast class. Since a proposition addresses a contrast, the bad because-claim is immediately disqualified. Spelled out in more detail, it is false that the flagpole is 30ft high because the shadow is 40ft long, because the proposition that the shadow is 40ft long does not address the contrast between the flagpole's being 30ft high and the flagpole's being any other height.

The second theory is van Fraassen's. On this view, a proposition addresses a contrast just in case what the proposition describes is properly probabilistically correlated with the elements of the contrast.[37] Crucially, this kind of probabilistic correlation does not distinguish between the two because-claims exhibiting the asymmetry of explanation, because *that* is symmetrical. So the proposition that the shadow is 40ft long addresses the contrast between the flagpole's being 30ft and being any other height. The second theory is therefore supplemented with

[36]For an extended treatment about how to spell out being causally relevant to a contrast, see Lipton (1991a,b).

[37]For some remarks on how proper probabilistic correlation is spelled out, see Van Fraassen (1980, pp. 146-51).

another degree of context-dependence which, on particular occasions of use, rules out the proposition that the shadow is 40ft long.

Let me make a couple of points about this way of setting up the choice of theories. Ordinarily, we want to choose between theories that differ as little as possible, ideally in just one respect. That way, we can hope to pinpoint where a theory whose predictions are not borne out goes wrong. The choice of theories I am presenting is not like this, since the two theories differ in two important respects. But that is unavoidable. It is a result of the fact that hypotheses about what addressing a contrast is interact with hypotheses about whether because-claims are sensitive to a relevance relation. If we do not want to impose a causal condition on addressing a contrast while maintaining metaphysical innocence, as well, we are forced to countenance a largely unconstrained relevance relation. That is why the competing theories differ both in their account of addressing a contrast *and* in whether they countenance sensitivity to a relevance relation.

Since the theories differ in two respects, one might hope to differentiate them according to two kinds of predictions they make. One of this I will pursue: the different predictions they make about whether because-claims exhibit dependence on a relevance relation. The other one concerns the necessity of causation for the truth of a because-claim. The causal theory predicts that every because-claim entails a suitable causal claim. So one could argue against the causal theory by exhibiting because-claims that are true even though the corresponding causal claim is false. That is the question usually addressed in the guise of "is there non-causal explanation?" *That* topic has seen a lot of discussion, and the debate seems to be deadlocked. That is why I won't discuss this second disagreement.

The problem is to find reasons to accept van Fraassen's theory over its competitor, and specifically his recognition of dependence on a relevance relation. These reasons need to be independent of his desire for metaphysical innocence. The best way is to find examples that would force even the causal theorist to recognize dependence on a contextually given relevance relation. Three examples of van Fraassen's are aimed at doing that. I'll discuss these in turn.

We have already seen the first example, repeated here.

(8) Why did the conductor warp rather than retain its shape?

a. The conductor warped rather than retain its shape because a very strong current went through it.

b. The conductor warped rather than retain its shape because it was in the Earth's magnetic field.

Since the two because-claims (8a) and (8b) explicitly mention a contrast, whatever difference in acceptability there is between them cannot arise from a difference in the contrast addressed. Further, both are plausibly interpreted as causal by the lights of theory we are considering, so if there is a difference, it cannot be due to one because-claim's satisfying a causal condition, the other's not. So if van Fraassen can show that there is a difference in truth-value between these two claims in a single context, then even the causal theory has to recognize a second, distinctive form of context-dependence of because-claims.

It should be uncontroversial that there are many contexts in which one of the because-claims is a better thing to say than the other. If we worry about how to prevent such warpings in the future, it is much better to be told that the conductor warped because a strong current went through it, not because it was in the Earth's magnetic field. That is not enough to support van Fraassen's desired conclusion: we need to show that, in this context, the because-claim about the field is *false*.

But it does not seem to be false. Even in a context where the because-claim about the current is much better, one can follow it up by adding the because-claim as addition information. That is, the discourse (20) is acceptable.

(20) A: Why did the conductor warp?

B: It warped because a strong current went through, and also because it was in the Earth's magnetic field.

However, the fact that this discourse is acceptable is not completely compelling by itself. It is possible that in the course of the assertion, the context changes. We have already seen the kind of evidence that helps us decide between these two alternatives. We need to check whether the context shifts in the course of evaluating the conjunction.

In the case of the conductor, we cannot resist the alleged shift in context in the way we can in the example of shape I discussed earlier in (17). We cannot felicitously reject the assertion that allegedly induces the shift in context, as (21) shows.

(21) A: Why did the conductor warp?

B: It warped because a strong current went through, and also because it was in the Earth's magnetic field.

A: # No, it didn't warp because of it was in the Earth's magnetic field.

The fact that *A*'s response is unacceptable strongly suggests that both because-claims are true. That means that we do not have reason to countenance a further kind of distinctive context-dependence.

Van Fraassen's other examples follow the same pattern.[38]

(22) a. I woke up because the milkman woke me up.

b. I woke up for no reason—I had no plans.

(23) a. Blood circulates through the body because the heart pumps blood through the arteries.

[38]The examples are taken from Van Fraassen (1980, pp. 141-2).

b. Blood circulates through the body because every part of the body tissue needs oxygen.[39]

As before, (22a) and (22b) happily go together. (22b) can be smoothly continued with "...I woke up because the milkman woke me." Likewise, "blood circulates through the body because the heart pumps it through the arteries and every part of the body tissue needs oxygen" is fine.

So far, I have only argued that the causal theorist need not countenance a further kind of context-dependence on a salient relevance-relation. That means that the causal theory is consistent with all of the evidence so far. Hence we have a clear judgment about the merits of the different theories: the causal theory is consistent with all of the evidence I have considered so far, while van Fraassen's theory is not. We should reject van Fraassen's theory.

6 Conclusion

I have argued that UNIFORMITY should be interpreted as a claim about the semantics of because-claims. That is the case, because many of the issues, such as the asymmetry of explanation, that traditionally have exercised philosophers of science concern the semantics of these claims. I have also argued that van Fraassen's

[39]This example is slightly modified from van Fraassen's text, where he considers "Blood circulates through the body in order to bring oxygen to every part of the body tissue." I have modified the example to ensure that we deal with a because-claim.

challenge to UNIFORMITY, thus construed, fails. Because-claims are not context-dependent in a way that threatens the philosophical project of giving an account of explanation.

References

Asher, N. and M. Morreau. "What Some Generic Sentences Mean". In G. N. Carlson and F. J. Pelletier, eds., *The Generic Book*, 300–339 (Chicago: University of Chicago Press, 1995).

Barr, W. F. "A Syntactic and Semantic Analysis of Idealizations in Science". *Philosophy of Science*, 38(2), (1971), 258–272.

—. "A Pragmatic Analysis of Idealizations in Physics". *Philosophy of Science*, 41(1), (1974), 48–64.

Boër, S. E. and W. G. Lycan. *Knowing Who* (Cambridge, MA: MIT Press, 1986).

Boolos, G. S. and R. C. Jeffrey. *Computability and Logic, Third Edition* (Cambridge, UK: Cambridge UP, 1989).

Boyd, R. N. "Realism, Anti-Foundationalism and the Enthusiasm for Natural Kinds". *Philosophical Studies*, (61), (1991), 127–148.

—. "Kinds, Complexity and Multiple Realization: Comments on Millikan's "Historical Kinds and the Special Sciences"". *Philosophical Studies*, 95, (1999), 67–98.

Brandon, R. N. "Does Biology have Laws? The Experimental Evidence". *Philosophy of Science*, 64, Supplement, (1997), S444–S457.

Bromberger, S. "Science and the Forms of Ignorance". In *On What We Know We Don't Know*, 112–127 (Chicago: The University of Chicago Press, 1992a).

—. "A Theory about the Theory of Theory and about the Theory of Theories". In *On What We Know We Don't Know*, 52–74 (Chicago: The University of Chicago Press, 1992b).

—. "Why-Questions". In *On What We Know We Don't Know*, 75–100 (Chicago: The University of Chicago Press, 1992c).

Cartwright, N. *How the Laws of Physics Lie* (Oxford: Clarendon Press, 1983).

—. "*Ceteris Paribus* Laws and Socio-Economic Machines". *Monist*, 78, (1995), 276–297.

—. *The Dappled World: A Study of the Boundaries of Science* (Cambridge, UK: Cambridge UP, 1999a).

—. "Models and the Limits of Theory: Quantum Hamiltonians and the BCS Model of Superconductivity". In M. Morrison and M. S. Morgan, eds., *Models as Mediators*, 241–281 (Cambridge, UK: Cambridge UP, 1999b).

Cartwright, N., T. Shomar, and M. Suárez. "The Toolbox of Science". In W. Herfel, W. Krajewksi, I. Niiniluoto, and R. Wójcicki, eds., *Theories and Models in Scientific Processes* (Amsterdam: Rodopi, 1995).

Dahl, O. "On Generics". In E. Keenan, ed., *Formal Semantics of Natural Language*, 99–111 (Cambridge: Cambridge UP, 1975).

Dray, W. *Laws and Explanations in History* (Oxford: Oxford UP, 1957).

Dupré, J. "Probabilistic Causality Emancipated". In P. French, J. Theodore Uehling, and H. Wettstein, eds., *Midwest Studies in Philosophy IX: Causality and Causal Theories*, 169–175 (Minneapolis: University of Minnesota Press, 1984).

Earman, J. and J. Roberts. "*Ceteris Paribus*, There is no Problem of Provisos". *Synthese*, (118), (1999), 439–478.

Earman, J., J. Roberts, and S. Smith. "*Ceteris paribus* Lost". *Erkenntnis*, 57(3), (2002), 281–301.

Farkas, D. F. and Y. Sugioka. "Restrictive *If/When* Clauses". *Linguistics and Philosophy*, (6), (1983), 225–258.

Fodor, J. A. "You can fool some of the People All of the Time, Everything Else Being Equal; Hedged Laws and Psychological Explanations". *Mind*, 100(1), (1991), 19–34.

—. "Making Mind Matter More". In *A Theory of Content and other Essays*, 137–159 (Cambridge: MIT Press, 1994).

French, S. and J. Ladyman. "Semantic Perspective on Idealization in Quantum Mechanics". In N. Shanks, ed., *Idealization IX: Idealization in Contemporary Physics*, 51–73 (Amsterdam: Rodopi, 1998).

Friedman, M. "Explanation and Scientific Understanding". *The Journal of Philosophy*, 71(1), (1974), 5–19.

Garfinkel, A. *Forms of Explanation: Rethinking the Questions in Social Theory* (New Haven: Yale UP, 1981).

Giere, R. N. *Explaining Science* (Chicago: University of Chicago Press, 1988).

—. *Science Without Laws* (Chicago: The University of Chicago Press, 1999).

—. "How Models are Used to Represent Reality". *Philosophy of Science*, 71(4), (2004), 742–752.

Glymour, C. "A Semantics and Methodology for Ceteris Paribus Hypotheses". *Erkenntnis*, 57(3), (2002), 395–405.

Goodman, N. *Fact, Fiction, and Forecast* (Cambridge, MA: Harvard UP, 1983).

Grant, I. S. and W. R. Phillips. *The Elements of Physics* (Oxford: Oxford UP, 2001).

Grice, H. "Logic and Conversation". In *Studies in the Ways of Words*, 22–40 (Cambridge: Harvard UP, 1991).

Griffiths, P. E. "Darwinism, Process Structuralism, and Natural Kinds". *Philosophy of Science*, 63, Supplement, (1996), S1–S9.

Grimes, T. R. "Explanation and the Poverty of Pragmatics". *Erkenntnis*, 27, (1987), 79–92.

Groenendijk, J. and M. Stokhof. "Questions". In J. van Bentham and A. ter Meulen, eds., *Handbook of Logic and Language*, 1055–1124 (Amsterdam: Elsevier and MIT Press, 1997).

Hacking, I. "A Tradition of Natural Kinds". *Philosophical Studies*, (61), (1991), 109–126.

Hall, N. "Causation and Ceteris Paribus Laws". *The Harvard Review of Philosophy*, 13(1), (2005), 80–99.

Hamblin, C. L. "Questions". *Austalasian Journal of Philosophy*, 36, (1958), 159–168.

—. "Questions in Montague English". *Foundations of Language*, 10, (1973), 41–53.

Haspelmath, M. *Understanding Morphology* (London: Arnold, 2002).

Hausman, D. M. *Causal Asymmetries* (Cambridge, UK: Cambridge UP, 1998).

Heim, I. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts, Amherst (1982).

Hempel, C. G. "Aspects of Scientific Explanation". In *Aspects of Scientific Explanation*, 331–496 (New York: The Free Press, 1965a).

—. "The Function of General Laws in History". In *Aspects of Scientific Explanation*, 231–243 (New York: The Free Press, 1965b).

—. "Typological Methods in the Natural and the Social Sciences". In *Aspects of Scientific Explanation*, 155–172 (New York: The Free Press, 1965c).

—. "Provisos: A Problem Concerning the Inferential Function of Scientific Theories". *Erkenntnis*, 28, (1988), 147–164.

Hempel, C. G. and P. Oppenheim. "Studies in the Logic of Explanation". In *Aspects of Scientific Explanation*, 245–295 (New York: The Free Press, 1965).

Hesse, M. *Models and Analogies in Science* (Notre Dame, IN: University of Notre Dame Press, 1966).

Hitchcock, C. R. "The Role of Contrast in Causal and Explanatory Claims". *Erkenntnis*, 107, (1996), 395–419.

Kaplan, D. "Demonstratives". In J. Almog, J. Perry, and H. Wettstein, eds., *Themes from Kaplan*, 481–614 (New York: Oxford UP, 1989).

Karttunen, L. "Syntax and Semantics of Questions". *Linguistics and Philosophy*, 1(1), (1977), 3–44.

Kincaid, H. "Global Arguments and Local Realism in the Social Sciences". *Philosophy of Science*, 67, Supplement(II), (2000), S667–S678.

Kitcher, P. "Explanatory Unification and the Causal Structure of the World". In P. Kitcher and W. C. Salmon, eds., *Minnesota Studies in the Philosophy of Science Vol. 13: Scientific Explanation*, 410–505 (Minneapolis: University of Minnesota Press, 1989).

—. "Explanatory Unification". In R. Boyd, P. Gasper, and J. Trout, eds., *The Philosophy of Science*, 329–347 (Cambridge, MA: MIT Press, 1991).

—. *The Advancement of Science* (New York: Oxford UP, 1993).

Kitcher, P. and W. Salmon. "Van Fraassen on Explanation". *The Journal of Philosophy*, 84(6), (1987), 315–330.

Krifka, M., F. J. Pelletier, G. N. Carlson, A. ter Meulen, G. Chierchia, and G. Link. "Genericity: An Introduction". In G. N. Carlson and F. J. Pelletier, eds., *The Generic Book*, 1–124 (Chicago: University of Chicago Press, 1995).

Kuhn, T. S. *The Structure of Scientific Revolutions* (Chicago: University of Chicago Press, 1962).

Kukla, A. "Scientific Realism, Scientific Practice, and the Natural Ontological Attitude". *The British Journal for the Philosophy of Science*, 45(4), (1994), 955–975.

Lange, M. "Are there Natural Laws Concerning Particular Biological Species". *The Journal of Philosophy*, 92(8), (1995), 430–451.

—. *Natural Laws in Scientific Practice* (Oxford: Oxford UP, 2000).

—. "The Autonomy of Functional Biology: A Reply to Rosenberg". *Biology and Philosophy*, 19, (2004), 93–109.

Lawler, J. *Studies in English Generics* (Ann Arbor: University of Michigan Press, 1973).

Laymon, R. "Idealization, Explanation, and Confirmation". *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1980(1), (1980), 336–350.

—. "Scientific Realism and the Hierarchical Counterfactual Path from Data to Theory". *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Assoication*, 1982(1), (1982), 107–121.

—. "Cartwright and the Lying Laws of Physics". *The Journal of Philosophy*, 86(7), (1989), 353–372.

Lewis, D. K. "Scorekeeping in a Language Game". In *Philosophical Papers, Vol. 1*, 233–249 (New York: Oxford UP, 1983).

—. "Causal Explanation". In *Philosophical Papers, Vol. 2*, 214–240 (New York: Oxford UP, 1986).

Lipton, P. "Contrastive Explanation and Causal Triangulation". *Philosophy of Science*, 58(4), (1991a), 687–697.

—. *Inference to the Best Explanation* (London: Routledge, 1991b).

Mayr, E. "Typological versus Population Thinking". In E. Sober, ed., *Conceptual Issues in Evolutionary Biology, Second Edition*, 157–160 (Cambridge, MA: MIT Press, 1994).

Mill, J. S. *A System of Logic: Ratiocinative and Inductive* (Honolulu: University Press of the Pacific, 2002).

Miller, R. W. *Fact and Method* (Princeton: Princeton UP, 1987).

—. "Half-Naturalized Social Kinds". *Philosophy of Science*, 67, Supplement(II), (2000), S640–S652.

Millikan, R. G. "Historical Kinds and the "Special Sciences"". *Philosophical Studies*, 95, (1999), 45–65.

Morreau, M. "Other Things Being Equal". *Philosophical Studies*, 96, (1999), 163–182.

Morrison, M. "Unification, Realism, and Inference". *The British Journal for the Philosophy of Science*, 41(3), (1990), 305–322.

—. "Models as Autonomous Agents". In M. Morrison and M. S. Morgan, eds., *Models as Mediators*, 38–65 (Cambridge, UK: Cambridge UP, 1999).

Morrison, M. and M. S. Morgan. "Models as Mediating Instruments". In M. Morrison and M. S. Morgan, eds., *Models as Mediators*, 10–37 (Cambridge, UK: Cambridge UP, 1999).

Mott, P. "Fodor and Ceteris Paribus Laws". *Mind*, 101(402), (1992), 335–346.

Nickel, B. "Semantics for Characterizing Sentences". In J. Gajewski, V. Hacquard, B. Nickel, and S. Yalcin, eds., *Recent Work on Modality* (Cambridge, MA: MIT Press, forthcoming).

Niiniluoto, I. "Verisimilitude: The Third Period". *The British Journal for the Philosophy of Science*, 49(1), (1998), 1–29.

Nowak, I. and L. Nowak. ""Model(s)" and "Experiment(s)" as Homogeneous Families of Notions". In N. Shanks, ed., *Idealization IX: Idealization in Contemporary Physics*, 35–50 (Amsterdam: Rodopi, 1998).

Oddie, G. "Truthlikeness". In E. N. Zalta, ed., *The Stanford Encyclopedia of Philosophy* (Fall 2001).
URL: *http://plato.stanford.edu/archives/fall2001/entries/truthlikeness/*

Okasha, S. "Darwinian Metaphysics: Species and the Question of Essentialism". *Synthese*, (131), (2002), 191–213.

Pietroski, P. M. and G. Rey. "When Other Things Aren't Equal: Saving Ceteris Paribus Laws from Vacuity". *The British Journal for the Philosophy of Science*, 46(1), (1995), 81–110.

Pinker, S. *Words and Rules: The Ingredients of Language* (New York: Perennial, 2000).

Popper, K. *Conjectures and Refutations: The Growth of Scientific Knowledge* (London: Hutchinson, 1963).

Putnam, H. "What Theories Are Not". In R. E. Grandy, ed., *Theories and Observation in Science*, 111–123 (Englewood Cliffs, NJ: Prentice Hall, 1973).

Quine, W. "Natural Kinds". In *Ontological Relativity and Other Essays*, 114–138 (New York: Columbia UP, 1969).

Rappaport, S. *Models and Reality in Economics* (Cheltenham, UK: Edward Elgar, 1998).

Rosenberg, A. "How is Biological Explanation Possible?" *British Philosophy of Science*, 52, (2001a), 735–760.

—. "Reductionism in a Historical Science". *Philosophy of Science*, 68, (2001b), 135–163.

Rudner, R. *Philosophy of Social Science* (Englewood Cliffs: Prentice-Hall, 1966).

Salmon, W. C. *Scientific Explanation and the Causal Structure of the World* (Princeton, New Jersey: Princeton UP, 1984).

—. "Four Decades of Scientific Explanation". In P. Kitcher and W. C. Salmon, eds., *Minnesota Studies in the Philosophy of Science Vol. 13: Scientific Explanation*, 3–219 (Minneapolis: University of Minnesota Press, 1989).

—. "Explanatory Asymmetry: A Letter to Professor Adolf Grünbaum from His Friend and Colleague". In *Causality and Explanation*, 164–177 (Oxford: Oxford UP, 1998).

Sandborg, D. "Mathematical Explanation and the Theory of Why-Questions". *The British Journal for the Philosophy of Science*, 49(4), (1998), 603–624.

Schiffer, S. "Ceteris Paribus Laws". *Mind*, 100(1), (1991), 1–17.

Schurz, G. "Pietroski and Rey on Ceteris Paribus Laws". *British Journal for the Philosophy of Science*, (52), (2001), 359–370.

—. "*Ceteris paribus* Laws: Classification and Deconstruction". *Erkenntnis*, 57(3), (2002), 351–372.

Schwartz, R. J. "Idealizations and Approximations in Physics". *Philosophy of Science*, 45(4), (1978), 595–603.

Shapere, D. "Notes towards a Post-Positivistic Interpretation of Science". In P. Achinstein and S. Barker, eds., *The Legacy of Logical Positivism*, 115–60 (Baltimore: Johns Hopkins Press, 1969).

Sklar, L. *Theory and Truth: Philosophical Critique within Foundational Science* (Oxford: Oxford UP, 2000).

Smith, S. "Violated Laws, *Ceteris Paribus* Clauses, and Capacities". *Synthese*, (130), (2002), 235–264.

Sober, E. "Evolution, Population Thinking, and Essentialism". In E. Sober, ed., *Conceptual Issues in Evolutionary Biology, Second Edition*, 161–189 (Cambridge, MA: MIT Press, 1994).

Stanley, J. C. "Context and Logical Form". *Linguistics and Philosophy*, 23(4), (2000), 319–434.

—. "Making it Articulated". *Mind and Language*, 17(1 & 2), (2002), 149–168.

Strevens, M. "The Causal and Unification Approaches to Explanation Unified—Causally". *Noûs*, 38(1), (2004), 154–176.

Suárez, M. "The Role of Models in the Application of Scientific Theories: Epistemological Implications". In M. Morrison and M. S. Morgan, eds., *Models as Mediators*, 168–196 (Cambridge, UK: Cambridge UP, 1999).

Suppe, F. "Introduction to *The Structure of Scientific Theories*". In F. Suppe, ed., *The Structure of Scientific Theories*, 3–232 (Urbana: University of Illinois Press, 1977).

Suppes, P. "What is a Scientific Theory?" In S. Morgenbesser, ed., *Philosophy of Science Today*, 55–67 (New York: Basic Books, 1967).

Teller, P. "Twilight of the Perfect Model Model". *Erkenntnis*, 55(3), (2001), 393–415.

Temple, D. "The Contrast Theory of Why-Questions". *Philosophy of Science*, 55(1), (1988), 141–151.

Van Fraassen, B. *The Scientific Image* (Oxford: Oxford UP, 1980).

Weston, T. "Approximate Truth and Scientific Realism". *Philosophy of Science*, 59, (1992), 53–74.

Woodward, J. "There is no such Thing as a *ceteris paribus* Law". *Erkenntnis*, 57(3), (2002), 303–328.

—. *Making Things Happen* (Oxford: Oxford UP, 2003).