How General Do Theories of Explanation Have to Be?

Citation

Nickel, Bernhard. 2010. How general do theories of explanation have to be? *Nous* 44(2): 305-328.

Published Version

doi:10.1111/j.1468-0068.2010.00741.x

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:4692275>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

How General Do Theories of Explanation Have To Be?*

Bernhard Nickel · Harvard University

June 18, 2008
Forthcoming in *Noûs*

Abstract

Theories of explanation seek to tell us what distinctively explanatory information is. The most ambitious ones, such as the DN-account, seek to tell us what an explanation is, *tout court*. Less ambitious ones, such as causal theories, restrict themselves to a particular domain of inquiry. The least ambitious theories constitute outright skepticism, holding that there is no reasonably unified phenomenon to give an account of. On these views, it is impossible to give any theories of explanation at all. I argue that both the less ambitious and outright skeptical varieties are committed to a certain context-sensitivity of our explanatory discourse. And though this discourse is almost certainly context-sensitive in *some* respects, it does not exhibit the context-sensitivity less than fully ambitious theories are committed to. Therefore, all accounts that seek to restrict themselves in scope, including causal accounts of explanation, fail.

1 Introduction

When we begin philosophical theorizing, we are inclined to use the concept of an explanation in the course of elucidating various problems. We might distinguish different

*This paper grew out of work I originally did as part of my dissertation at MIT. I had wonderful teachers there, and I owe a great debt of gratitude to Ned Hall, Bob Stalnaker, Judy Jarvis Thomson, and Steve Yablo. While I was there, and in the time since I left, I have been helped a lot by Paul Pietroski and Jason Stanley. Eric Swanson was my office mate for several years, and discussed many versions of the arguments in this paper. Agustín Rayo forced me to think through the possibility of irresistible accommodation. The paper went through one final major overhaul in response to comments made by an anonymous referee for this journal.

theories of the mind, such as behaviorism, physicalism, or functionalism, in terms of the explanations they sanction. We might characterize the debate between realists and their opponents in the philosophy of science in terms of explanations that are available. And philosophical inquiry itself often aims self-consciously at explanations of the phenomena that constitute its subject matter: language, the mind, the difference between right and wrong. Theories of explanation have been central to the philosophical enterprise.

In the usual case, a theory of explanation seeks to tell us what an explanation is by telling us what distinctively explanatory information is.¹ Deductive-nomological, causal, and unificationist accounts all fit into this pattern. But such theories may be more or less ambitious. They may seek to say what makes information explanatory *tout court*, or they may seek to do so only for a limited domain, such as natural science. They are also distinguished by whether they give necessary conditions, sufficient ones, or both. The DN-account of explanation is as ambitious as such theories can be. It gives necessary and sufficient conditions, and it makes no distinctions of subject matter: whether we want to explain why an event occurred, why a law holds, or why an action is right, we always need to produce a valid argument that satisfies the constraints of the account.²

However, the DN-account faces serious difficulties, and the most dominant approach to explanation to replace it has also dialed back its ambitions. This is the causal account of explanation, on which information about an event is explanatory only if it is causal information about that event. Its proponents are only concerned with giving a necessary condition on a bit of information to be explanatory, and the account is only meant to apply in domains in which causal relations obtain. Explanations of laws fit only awkwardly into this paradigm, explanations in mathematics and ethics not at all.

Restrained ambition need not be a bad thing. Perhaps the insight of causal theories of explanation is just that no general account of explanation is to be had. There is no single phenomenon to give an account of, and to look for more is a mistake. Perhaps we have simply found out what the natural kinds of explanation are. Thus causal theories of explanation paradigmatically endorse the claim I'll call

[MODERATION] Different domains of inquiry require different modes of explanation.

¹A notable exception is Achinstein (1983), who analyzes explanations fundamentally in terms of a distinctive kind of speech act that is not equivalent to the expression of a distinctive kind of information.

²The *locus classicus* for the DN-account is Hempel (1965); one of its most sophisticated developments is in Railton (1978, 1981). Classic unificationist accounts are Friedman (1974) and Kitcher (1989, 1991). Some causal accounts are given by, for example, Hausman (1998), Lewis (1986), Salmon (1984, 1989), Strevens (2004), and Woodward (2003).

This is a moderate position insofar as it embodies moderation in the ambition of theories of explanation. The crucial aspect of MODERATION is the commitment that across domains, we can give no substantive account of what makes something an explanation. For example, the moderate might give a very substantial account of explanation in the sciences, such as a causal account, and she might also give a substantive account of explanation in mathematics, perhaps in terms of an account of what makes a proof explanatory or illuminating.³ But according to a moderate, there is no substantive account of what makes explanations in natural science and explanations in mathematics *explanations*.⁴

At this point, we can ask just what counts as a domain for the purposes of interpreting MODERATION. For example, it may turn out that different scientific disciplines employ sufficiently different modes of explanation so that no substantive generalizations even within natural science are available. Thus, moderate views can be arranged along a spectrum, at one extreme of which is a view according to which there are only very few modes of explanation—perhaps as few as three or four, one each of science, ethics, and mathematics. At the other extreme stands a view I’ll call

[SKEPTICISM] There is no such thing as distinctively explanatory information. Hence, there can be no theory of explanation.

The basic thought of SKEPTICISM is that what an investigator will accept as an explanation is subject to a host of factors. These factors are so diverse, and their influence so pervasive, that they prevent us from identifying a common core of what makes something an explanation in even the smallest of domains, except in the trivial case where we address a single phenomenon on a particular occasion. This is the view propounded by Van Fraassen (1980, chp. 5).⁵

In this paper, I oppose all forms of moderation. I argue that they all fail to square with our explanatory discourse, specifically our practice of asking and answering why-questions.⁶ I will argue that there are substantive constraints on acceptable answers to

³As for example Sandborg (1998) and Steiner (1978a) propose.

⁴An illuminating analogy may be with Tarski’s definition of truth-in-L, which did not (nor was designed to) provide any substantive account of what makes truth-in-L and truth-in-L’ both *truth*.

⁵This paper originates in trying to think through the arguments in that very rich text. However, I do not engage van Fraassen’s text directly, largely because I think that the positions I discuss—and which are similar to suggestions he makes—are the best options for a skeptic about explanation. I thus hope to bypass lengthy exegetical issues and go directly to what I think is the greatest challenge we can extract from that text.

⁶At times, writers simply identify explanations with answers to why-questions (see, e.g., Van Fraassen (1980, p. 134)). As I point out at the end of the paper (p. 24), that cannot be right. Asking and answering why-questions is only part of our explanatory discourse, not all of it. However, for the main part of the paper, I will speak as if this identification was unobjectionable.

why-questions, and that these constraints are invariant across all domains. In somewhat telegraphic form, which I will expand on below, I will argue that asymmetries of explanation are due to context-invariant constraints.

After some preliminary remarks to set the stage in section 2, I explain which aspect of our explanatory discourse I focus on, as well as how my argument relates to explanatory asymmetries in section 3. This will also make precise my claim that the context-invariant constraints I argue for are “substantive.” Section 4 argues that moderates and skeptics are both committed to a particular semantic theory. Sections 5 and 6 contain a two-step argument against my main target, the semantic commitments of MODERATION. The first step in that argument, presented in section 5, also serves as an argument against SKEPTICISM by arguing that its semantic commitments are untenable.

2 The Implications of SKEPTICISM and MODERATION

All parties to the debate—the skeptic, the moderate, and I—agree that on a particular occasion, we accept only certain, quite specific bits of information as explanatory. The basic disagreement is over how much of that selectivity reflects general constraints on explanatory information, and how much of it reflects specific features of the situation that may vary from situation to situation, constraints imposed by interests that we bring to the table that go beyond simply our interest in an explanation.

Skepticism

The skeptic’s position is the most extreme. On her view, there are no invariant constraints on what can count as explanatory information for a given explanandum (save that the explanation be true): for any pair of true propositions, there is a context in which one counts as an explanation of the other. Here are three consequences of this view.

Traditionally, *understanding* a phenomenon is taken to be a distinctive epistemic achievement, more than merely being able to predict it. Understanding a phenomenon is having an explanation of it. An account of understanding might then be a major component of a theory of science, since one might hold that science aims, *inter alia*, at understanding the phenomena it investigates. If SKEPTICISM is correct, the special status of understanding is a chimera, because there is no distinctive explanatory information. Instead, it is simply ordinary information, found in a context we happen to care about it. We might call such episodes of finding information “understanding,” but there

would be nothing that unifies them, and hence nothing that could serve as a general theory of science.

Another example. According to many theorists, inference to the best explanation (IBE) can play a major role in solving the problem of induction.⁷ For any given set of data, there are infinitely many hypotheses that are compatible with the data but diverge on their predictions regarding as yet unobserved phenomena. Among these hypotheses, only some are supported by the data. The problem is to say which. According to IBE, the supported hypotheses would, if true, offer explanations of the observed data, the rest would not. If SKEPTICISM is true, IBE is undermined. For each of the competing hypotheses, there is some context in which it counts as a potential explanation of the data. Being a potential explanation therefore does not distinguish any hypothesis from any other.

Third example. A successful reduction plausibly requires more than supervenience, since logical space leaves room for non-reductive supervenience relations. On one popular proposal, a reduction requires that we be able to *explain* the phenomena to be reduced by appeal to the reduction base.⁸ If SKEPTICISM is correct, we have to ask why being able to provide a kind of information we count as explanatory in some contexts but not in others indicates that a metaphysically important relation obtains.

This skeptical position is certainly very extreme and may strike some readers as so implausible as to not be worth discussing, so let me say why I discuss it nonetheless. First, this extreme version has the easiest time in arguing for the radical conclusions I just mentioned. It may be possible to establish similar conclusions based on slightly weaker premises, saying roughly that the context-dependence of our judgments of explanatoriness is pervasive and hard to factor out. But the arguments become much murkier once we draw fine distinctions between degrees of context-sensitivity.

Second, if I succeed in arguing against MODERATION, then *all* forms of it are mistaken, including the extreme form I here single out under the heading of SKEPTICISM, as well as more attenuated versions that still aim at the skeptical conclusions. So even though I do not discuss these intermediate versions directly, they will nonetheless be caught up in the web of my argument.

Third, this particular version of skepticism has been historically important, since it is the view advocated by van Fraassen in his (1980). One of my aims in discussing it is to show just how strong a claim about the context-dependence of our explanatory

⁷See, e.g., Harman (1965) and Lipton (1991a,b).

⁸See, e.g., Horgan (1993).

discourse a position like his is committed to.

Fourth and finally, I want to use SKEPTICISM as an illustration of the first stage of my argument, and for that reason, it will be expedient to formulate it as crassly as possible. This will allow me to pursue the main lines of objections and replies to my arguments in a relatively simple setting. Nonetheless, this should not distract from the fact that my real quarry is MODERATION.⁹

Moderation

According to the moderate, there is no explanation *per se*, there is only explanation in some domain, such as perhaps the natural sciences, ethics, or mathematics. In each case, the demands a piece of information has to meet in order to count as explanatory differ. Which demands are relevant depends on the context, specifically, on the kind of investigation we are conducting.¹⁰ While moderation is more palatable than skepticism, it has none of its philosophically powerful consequences. Even if the moderate is right, there can be successful treatments of understanding, inference to the best explanations, and reduction. It just turns out that each of these accounts will be subject-matter specific.

The interest of moderation is primarily methodological. First, it would be an important guide to our theorizing about explanation, since it would place boundaries on the scope of any theory of explanation. And second, it would require us to be more cautious in exporting techniques from one domain of inquiry to another. For example, it is commonplace to argue in philosophy by drawing an inference to the best explanation, pointing to the fact that this is acceptable in the natural sciences. But if moderation is true, explanation in the natural sciences may well make IBE a good kind of inference there, while explanation in philosophy fails to underwrite it.

⁹This should also make clear that my argument is different from the anti-context-sensitivity arguments offered by Cappelen and Lepore (2005). They argue roughly as follows. If we accept context-sensitivity in certain relatively plausible cases, we are forced to accept context-sensitivity in completely unacceptable ones, as well. By contrast, I do not argue that MODERATION leads to SKEPTICISM and then argue by *modus tollens*. Rather, I argue that MODERATION fails by its own lights.

¹⁰This kind of position may be very sympathetic to the views of Nancy Cartwright as put forth in her (1999). She argues there that we should not think of scientific inquiry as even aiming at a single, unified account of the world. Rather, each discipline, or perhaps sub-discipline, investigates its own domain, in the process often creating the conditions required for successful inquiry, albeit in a temporally and spatially bounded area, such as a lab or a test bed. It would be a natural extension of this picture of science that different sub-disciplines not only use their domain-specific modes of investigation, but also their domain-specific modes of explanation.

3 Focusing on Because-Claims

I said that a theory of explanation has to square with our explanatory discourse, specifically with our practice of asking and answering why-questions. This general slogan leaves open exactly how explanations are related to this discourse. An explanation could just be a true because-claim, a sentence of the form *p because q*. Alternatively, an explanation could be a true and direct answer to a why-question, that is to say, a true because-claim that counts as a direct answer to the question asked (or assumed in the context).^{11,12}

Answerhood conditions enter the picture this way. In the course of giving semantics for ordinary indicative sentences, we can associate them with Davidsonian truth-conditions or propositions as their semantic values. Such propositions may be sets of possible worlds, Russellian complexes, or Fregean senses. All of these semantic values bring with them the commitment that the sentences whose semantic values they are can be true or false. But this strategy is obviously out of place for interrogative sentences, which are not true or false. Instead, much current work on interrogatives originates with the idea that their semantic values are answerhood conditions, rather than truth-conditions, i.e., conditions a proposition has to satisfy in order to count as an answer to the question expressed by the interrogative.¹³

The two options about how explanations are related to our explanatory discourse are substantially different, since answerhood conditions are usually more context-sensitive than truth-conditions. To see this, consider the somewhat simpler case of who-questions. Even though it may be true that Bill Clinton is the guy at the bar, the sentence “Bill Clinton is the guy at the bar” could fail to also express a direct answer to “Who is Bill Clinton?” on a particular occasion, since the context may impose some restriction on the possible answers, ruling out that response but allowing “Bill Clinton is the 42nd President.” In another context, the answerhood conditions may be different. All the

¹¹This latter option is the one taken by Van Fraassen (1980, p. 134).

¹²The restriction to *direct* answers is important but hard to make precise. The dialogue in (A) can occur quite naturally.

- (A) i. Why are the tides always high on opposite sides of the Earth?
 ii. Newton is really important here.

The person giving the response might be trying to jog the questioner’s memory with a hint that is supposed to point to gravity, but in so doing, she hasn’t given an explanation, and that is because the response is not a direct answer. For more on directness and some of its attendant complications, see Kadmon (2001, p. 13).

¹³Good starting points for these semantic approaches to questions are Groenendijk and Stokhof (1997) and Hamblin (1958, 1973).

while, both sentences remain true throughout.¹⁴

However, the argument I wish to mount does not depend on whether explanations are just true because-claims or true, direct answers to why-questions. I will argue that there are substantive, context-invariant conditions a because-claim has to satisfy in order to be true. If an explanation just is a true because-claim, then it follows immediately that there are context-invariant conditions a bit of information has to satisfy in order to be explanatory information. Moreover, truth-conditions are inherited by the conditions a because-claim has to satisfy in order to be a true and direct answer to a salient why-question, so on that view of explanation, it follows once again that there are substantive, context-invariant constraints on explanatory information.¹⁵

Asymmetries of Explanation

Imagine that we are doing astronomy and thinking about the moon. In that case, we are happy to accept (1a), but we reject (1b).

- (1) a. The moon appears there because it was at location *l* earlier.
- b. The moon appears there because it will be at location *l'* later.

Let us stipulate that all of the component sentences are true—the moon really appears there, it really was at *l*, and it really will be at *l'*. This example exhibits what is usually called an asymmetry of explanation, and it is important because of the role it has

¹⁴For an elaboration of this view about who-questions, see Boër and Lycan (1986). Their view that who-questions impose context-sensitive answerhood conditions has been opposed by Braun (2006).

¹⁵Many views about because-claims and their semantics are compatible with my argument. However, I am committed to a very general hypothesis about the relationship between language and the information conveyed by uses of language. This hypothesis has been pervasive in philosophy of language, at least since the pioneering work of Montague (2002). That is the hypothesis that at least some of the information conveyed is semantically encoded by the language. In my particular case, I am committed to the hypothesis that the explanatory information conveyed by a because-claim (in context) is semantically expressed by that claim (in context). If this was not the case, my contention that a theory of explanation has to square with our explanatory discourse would be untenable.

Here's why. My argument is roughly that because there are context-invariant constraints on the interpretation of because-claims, there are context-invariant constraints on what counts as explanatory information. But here is a way for that inference to fail. Imagine that what explanatory information is conveyed by an assertion of a because-claim is at best only lightly constrained by the interpretation of the sentence asserted. Instead, what information is conveyed is largely a matter of how the interpretation of the sentence interacts with other aspects of the context, via Gricean and other mechanisms. In that case, even if the semantic interpretation is context-invariant across domains, since that interpretation hardly constrains the information conveyed, it may still be true that the information conveyed varies greatly across domains.

Many researchers in the field share my commitments regarding the relation between language and information, at least for this example. However, this commitment has come under attack, e.g., by Chomsky (2000) and Pietroski (2005). The most lively way that this debate has been pursued has taken the form of asking whether the context-sensitivity of the information conveyed by a use of language is plausibly mirrored by the context-sensitivity of the semantic interpretation of the sentences used to convey that information. Thus, the very large and sophisticated literature on context-sensitivity and implicature is relevant to this debate.

played in the downfall in the classical DN-account of explanation, as well as in the rise of its dominant competitors. On the DN-account, we can explain why an event occurred by citing laws and statements about contingent states or events that jointly entail a statement describing the explanandum.¹⁶ Unfortunately for the DN-account, both (1a) and (1b) should count as giving explanations by this criterion. In both cases, the present position of the moon can be deduced from the laws of planetary motion and the information cited.

Many theorists who endorse a causal account of explanation point to the asymmetry as crucial evidence in favor of their view, since they can readily diagnose the problem with (1b). It does not give causal information about the present position of the moon, and is thus disqualified from giving explanatory information, no matter what else is true.

I, too, will be focusing on the difference between (1a) and (1b), and I will argue that the difference between them is due to a context-invariant condition on explanatory information. Considering the centrality of explanatory asymmetries in adjudicating between theories of explanation, this conclusion amounts to saying that a *substantive* constraint on explanation is context-invariant.¹⁷

Since I said that I will be arguing for a context-invariant component to the truth-conditions of because-claims, my first order of business is to argue that the asymmetry in the acceptability of my target sentences (1a) and (1b) is due to a difference in truth-

¹⁶It's also required that all of the premises in this derivation are essential, an attempt to exclude intuitively irrelevant information from an explanation. This condition doesn't matter here, so I ignore it in the text.

¹⁷There is another class of cases that also goes under the heading of an asymmetry. An event e_1 can be predicted by citing certain laws along with some contingent event e_2 . The laws are chosen to be symmetrical, meaning that we could also predict the occurrence of e_2 by citing these same laws and the occurrence of e_1 . But intuitively, only the first is an explanation. An example of this second type of asymmetry is the famous flagpole case. Imagine that a flagpole is standing in an otherwise empty field on a sunny day, casting a shadow on the ground. Let's assume that the shadow is 40ft long, the flagpole 30ft high. In that case, (Bi) is judged as giving an explanation, (Bii) as not.

- (B) i. The shadow is 40ft long because the flagpole is 30ft high.
- ii. The flagpole is 30ft high because the shadow is 40ft long.

Given the laws of trigonometry and assumptions about the paths of the lightrays, we can derive the length of the shadow from the position of the sun and the height of the flagpole, and we can equally derive the height of the flagpole from the position of the sun and the length of the shadow. So once again, the DN-account cannot distinguish between two sentences that obviously differ with respect to whether explanatory information is being provided.

This was, in fact, the first asymmetry to be noticed. The author usually credited with the observation is Bromberger (1992b), who has examples with a similar structure.

Causal theorists diagnose the problem here the same way as they do the problem raised by (1a) and (1b) and give the same solution. I do not discuss the example in the main text because my argument focuses on examples in which the explanandum remains constant, as in (1a) and (1b), which is not the case in (Bi) and (Bii). However, if as seems plausible, the very same constraint is at work in (1) and (B), then my conclusion carries over to these cases, as well.

value. If that is right, it follows that the unacceptable (1b) is unacceptable because it is false.

A very important piece of evidence for the falsity of (1b) comes from negating it. If it was true but inappropriate, negating it would yield a false sentence, which would therefore also be inappropriate. However, if (1b) is actually *false*, negating it yields a true sentence. So consider the negation of (1b), (2).

(2) It is not the case that the Moon appears there because it will be at location l' later.

This sentence is completely acceptable, which means that it must be true. Therefore, the negated sentence (1b) must be false.

Here is a subsidiary argument. Not only do we not want to say (1b), we do not believe it, either. But all of the features of a context that make a sentence unacceptable in spite of being true do not carry over to belief. For example, a sentence may be unacceptable to say because it is uninformative or insulting to the audience. Neither of these conditions prevents a speaker from believing that the sentence is true. So (1b) must be false.

4 Semantics For Because-Claims for Skeptics and Moderates

In the course of this section, I'll develop a set of truth-conditions for because-claims, including claims about their context-sensitivity, that I think are the best bet for both the moderate and the skeptic. They are the best bet because they include the most plausible hypotheses about the context-sensitivity of because-claims, as well as being the weakest such hypotheses that yield the desired results.

I'll start with the assumption that the components of a because-claim must themselves be true in order for the because-claim to be true. So we can begin thus.

[TRUTH (PROVISIONAL)] p because q is true with respect to a context C iff

- (i) $\lceil p \rceil$ is true with respect to C , and
- (ii) $\lceil q \rceil$ is true with respect to C .

I've formulated TRUTH this way to take account of the fact that, in addition to the distinctive context-sensitivity of because-claims that is my focus in this paper, the component sentences may themselves contain context-sensitive expressions, such as tense.

We now need to add further conditions, since these provisional truth-conditions make ‘because’ equivalent to ‘and.’

Contrast and Relevance

Let’s begin with the most plausible and widely accepted hypothesis about the context-sensitivity of because-claims, sensitivity to a contrast. To introduce it, consider the biblical story of Adam and the apple. Given this background, consider now (3a) and (3b).

- (3) a. Adam ate the apple because he was hungry.
- b. Adam ate the apple because Eve offered it to him.

We can hear either of these as excellent claims about the situation or as claims that leave something to be desired. To bring this out, consider two variations on (3a), each adding a bit of material.

- (4) 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.

Given the biblical story as background, it seems clear that (4a) is true, (4b) false.¹⁸ Let’s say that the ‘rather than’ phrase introduces a contrast. Our sense that (3a) can sound better or worse in this situation can then be explained by saying that its interpretation is context-sensitive: it has to be evaluated with respect to a contextually given contrast, and depending on which contrast we use to interpret (3a), it either expresses something true or something false. More generally, we can say that because-claims are always evaluated with respect to a contrast, and that when the contrast is not mentioned explicitly, it has to be supplied by the context. A because-claim, then, is true only if the information offered addresses the contrast with respect to which the claim is evaluated. So given this terminology, (4a) and (4b) are evaluated with respect to the same contrast, but only (4a) addresses it, as well.¹⁹ That allows us to fill out our provisional truth-conditions a bit to yield.

¹⁸That is not to say that, if the facts about the situation were different, the proposition expressed by (4b) couldn’t be true. Just imagine that Adam was hungry and was told that the pear wouldn’t relieve his hunger.

¹⁹Lewis (1986) challenges this hypothesis about the interpretation of because-claims, suggesting that at least some of them are not evaluated with respect to any contrast. However, I will simply concede to my opponents that Lewis is wrong on this count.

[TRUTH (STILL PROVISIONAL)] *p because q* is true with respect to a context *C* iff

- (i) 'p' is true with respect to *C*, and
- (ii) 'q' is true with respect to *C*, and
- (iii) the proposition expressed by 'q' with respect to *C* addresses the contrast salient in *C*.

Recognizing this much context-sensitivity is not enough to account for the difference between (1a) and (1b) by itself without either giving a more substantive theory of what it takes to address a contrast, or hypothesizing another context-sensitive element in the interpretation of because-claims. In fact, I'll suggest that the former route isn't open to the moderate or the skeptic, so both should accept that there is another element of context-sensitivity.

One way of accounting for the difference between (1a) and (1b) is to give an account of addressing a contrast that has, as a consequence, that the information about the future position of the moon fails to address the salient contrast. One account that would fit the bill is the one outlined by Lewis (1986). A bit of information about the causal history of an event addresses the contrast with another event only if the information mentioned would not also be true of a causal history leading to the contrast event.²⁰ For example, Adam's eating the apple had as part of its causal history his being hungry. Had he eaten the pear, that event too would have had Adam's being hungry as part of its causal history, and for that reason, Adam's being hungry does not address the contrast between eating the apple and eating the pear. However, the alternative pear-eating presumably would not have included Eve's offering Adam a pear as part of its causal history, which is why Eve's offering Adam the apple does address the contrast.²¹

On this account of addressing a contrast, causality is built into the very nature of addressing a contrast, which in turn means that the same constraint, that explanatory information be causal information, is present in all cases. This directly conflicts with both MODERATION and SKEPTICISM, since it would mean that the same mode of explanation was applicable everywhere: causal explanation.

Rather, what is wanted is a way to impose some condition that is satisfied in (1a) but not (1b), and in a way that allows that condition to vary with context. The way to accomplish these goals is to embrace a relatively weak account of addressing a contrast, one that can be the same across all contexts, and then require that potentially

²⁰See Lewis (1986, pp. 229-30).

²¹This is just one way of adapting a causal account to deal with contrasts. For an alternative, see Lipton (1991a,b); Hitchcock (1996) gives an account in the setting of a probabilistic account of causation.

explanatory information address the contrast in a particular way, to be specified by context.

Such an account of addressing a contrast might be a probability-raising approach which focuses on rational subjective credence. On this view, a proposition addresses a contrast for a subject iff the subject would adjust her credences in some appropriate way, were she to add that proposition to a suitable subset of her beliefs. To remain with the example of Adam and the apple, perhaps she should increase her credence that Adam eats the apple while decreasing her credence that he juggles it, once she adds that Adam was hungry to some suitable subset of her beliefs. I mention a *suitable* subset of beliefs, rather than all of them, because if she is in a position to ask why Adam ate the apple rather than juggle it, she already believes that Adam ate the apple and did not juggle it, so adding the information that he was hungry to a credal state including *those* beliefs would not result in the proper change in credences. Doubtless, we also need to make the required changes in credence much more sophisticated.²²

This kind of account would no longer suffice to distinguish between (1a) and (1b). Rather, it is then up to the context to specify some further condition that the potentially explanatory information has to satisfy, in addition to addressing the contrast. The context might specify, for example, that the information must be causal information, but it might also be unifying information, or whatever the case may be.

On this strategy, we get the same effect on particular occasions as we do with the causal account I just mentioned. But neither a skeptic nor a moderate who takes this approach is committed to seeing the same causal constraint in all contexts. This gives us the final part of the truth-conditions for because-claims.

[TRUTH (FINAL)] p because q is true with respect to a context C iff

- (i) $\lceil p \rceil$ is true with respect to C , and
- (ii) $\lceil q \rceil$ is true with respect to C , and
- (iii) the proposition expressed by $\lceil q \rceil$ with respect to C addresses the contrast salient in C , and
- (iv) the proposition expressed by $\lceil q \rceil$ with respect to C satisfies the relevance relation salient in C .

The truth-conditions summarized in TRUTH are the ones both the skeptic and the moderate should accept. They embody the weakest hypothesis about the context-sensitivity

²²For more on these problems, see Salmon (1989, p. 56) and citations, as well as Van Fraassen (1980, chp. 5).

of because-claims that is strong enough to yield the conclusion that the asymmetry of explanation exhibited by (1a) and (1b) is due to a contextually variable feature of the interpretation of because-claims, and it makes use of the most plausible form of context-dependence.

Permissiveness

So far, moderates and skeptics agree on their semantic commitments. They part ways not over the aspects of the interpretation of because-claims that are context-sensitive, but over the semantic constraints on these context-sensitive aspects. Many words that are context-sensitive impose such semantic constraints. A clear example is the pronoun ‘she.’ When it is used deictically—as a demonstrative—it depends for its interpretation on the context. But there is also a semantic constraint, since ‘she’ can only be used to refer to a female. This constraint is part of the semantics of ‘she.’

The crucial claim the skeptic endorses, and which distinguishes her from the moderate, is that there are no semantic constraints on either the possible contrast or the possible relevance relations. Specifically, she endorses

[PERMISSIVENESS] Neither what it is to address a contrast, nor the possible relevance relations, are semantically constrained, nor do they constrain each other.

Let me make two remarks about PERMISSIVENESS. The first is that this is purely about the *semantic* constraints on the interpretation of because-claims. A skeptic can agree that we often look for a specific kind of information, perhaps even in most cases. But that fact, if it is a fact, is due to interests that go beyond our interest in an explanation.

Second, the last clause of PERMISSIVENESS, that contrast and relevance do not constrain each other, is important. It is what makes for the difference between a skeptic and a moderate. Without that clause, it’s possible that the contrast salient in a context determines the relevance relation with respect to which we evaluate the because-claim. For example, one might hold that when the contrast is between two events studied by natural science, the relevance relation is constrained to be causation, i.e., requires that explanatory information be causal information. This is simply the moderate position.²³

²³A skeptic can also agree that there are features that go into evaluating the aptness of a potential explanation, such as fit with other explanations, feasibility of testing the explanation in other experiments, or what have you. However, to the extent that even by the skeptic’s lights, these further factors do not influence the truth-conditions of because-claims, I ignore them here.

5 Attacking PERMISSIVENESS

My strategy is to focus on PERMISSIVENESS, and more specifically, on the relationship between the interpretation of contrast and relevance. I'll argue first that once we fix a salient contrast, there is no further contextual variability in the interpretation of because-claims. If that is right, then the final clause of PERMISSIVENESS fails, and hence SKEPTICISM fails as well.

However, that initial result is quite compatible with MODERATION. In light of the initial result, the moderate is simply committed to the claim that once we have fixed on a contrast, we have also fixed on a relevance relation. For all that, all contrasts between elements from a particular domain may require one relevance relation, while contrasts drawn from another domain require a different one. That would still bear out the moderate, since different domains would require different modes of explanation, cashed out now as different relevance relations.

I'll argue against the moderate's commitment in the second stage of my argument. I'll exhibit pairs of sentences in which the same contrast is addressed, but which come from different domains, such as mathematics and natural science, or ethics and natural science. These examples show that there is a single relevance relation at work in evaluating because-claims in all of these domains. And for that reason, causation cannot be the relevance relation, even in the domain of natural science, since that relation is not applicable in the other domains.

Contrast Fixes Relevance

In order to make my argument that contrast fixes relevance, I'll consider the best-case example the skeptic can point to. These are sets of sentences that differ in their acceptability in different contexts, even though they address the same contrast. I'll argue that in these examples, we have no evidence for the relevant context-sensitivity. Hence, we have no reason to believe in such context-sensitivity.

Here is one such case.²⁴ Imagine a situation in which an electrical conductor experiences a surge in current and warps. The conductor is on a space probe. Had the probe not been in the Earth's magnetic field, the conductor wouldn't have warped, nor would it have warped had the current not gone through it. Now consider two possible because-claims we could make about the situation, (5a) and (5b).

²⁴This example is modeled on the discussion in Van Fraassen (1980, pp. 141-2).

- (5) a. The conductor warped because a strong current went through it.
- b. The conductor warped because it was in the Earth's magnetic field.

Intuitively, there are contexts in which (5a) is a better thing to say than (5b) is, and there are contexts in which the opposite is true. For an example of the former, we might not know what happened inside the probe. An example of the latter is one in which we know about the current, but are puzzled nonetheless since space probes usually aren't in magnetic fields. This difference in acceptability cannot be due to a difference in salient contrasts, since the two responses address the same one.

More abstractly, here are the relevant features of the case. In one context, call it C_1 , (5a) is completely acceptable. That shows that in C_1 , (5a) is true. In that same context, we may assume, (5b) is less good. In another context, C_2 , (5b) is completely acceptable and hence true. In C_2 , (5a) is not as good. Let me also stipulate that the component sentences do not differ in their interpretation with respect to the two contexts C_1 and C_2 . The question is whether there is a context-sensitive element in the two because-claims whose interpretation changes across the two contexts C_1 and C_2 .

I'll argue that (5a) and (5b) are true with respect to C_1 and C_2 . In that case, we do not have any reason to believe that once we have fixed the salient contrast (as we have for (5a) and (5b)), there are any other contextually variable elements to the interpretation of these sentences.

To make that argument, consider the conjunction of the two because-claims (6).

- (6) The conductor warped because a strong current went through it and because it was in the Earth's magnetic field.

This conjunction is clearly acceptable, and hence true. There are two possible accounts for why (6) expresses a truth. On the first, the whole conjunction is evaluated with respect to the single context—call it C_3 —and hence one in which all the context-sensitive elements of (5a) and (5b) receive the same interpretation. If that is true, then the following argument holds. C_1 and C_3 agree on how to interpret (5a). C_2 and C_3 agree on how to interpret (5b). And C_3 is such that the context-sensitive features of (5a) and (5b) are interpreted the same way. Thus, C_1 and C_2 agree on how to interpret both (5a) and (5b). So if (5a) is true with respect to C_1 , it's also true with respect to C_2 , and if (5b) is true with respect to C_2 , it's also true with respect to C_1 .²⁵ This is the account I endorse.

²⁵One might worry about this argument. How does the fact that the conjunction (2) is true with respect to C_3 show that C_3 and C_1 agree on how to interpret the first conjunct? It is possible, after all, that the

But there is another account for why (6) is acceptable. If it is correct, the skeptic can maintain her position. On this second account, the context changes in the course of interpreting the conjunction. We interpret the first conjunct with respect to C_1 , and then the context shifts to another context C_2 . I need to rule out this alternative account.

If the alternative account is true, there must be some mechanism that forces the context to shift. We know that there are many possible mechanisms to accomplish this goal: we can shift the context by explicit announcement (“OK, now I’m interested in standing background conditions...”), or by changing the physical or intellectual environment in certain ways (e.g., by finishing one part of the report on what happened to the space probe). But none of these specific mechanisms need to operate in order for the conjunction (6) to be acceptable. Therefore, the assertion of the sentence itself must force the shift in context. This can happen via the mechanism Lewis (1983) calls *accommodation*.

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, we might be discussing basketball players. In that case, my utterance of “I am tall” is false. But I might then say: “I am tall. I tower over all the kids in my daughter’s day-care class.” Merely saying this can change the context so that what I said is true, even though with respect to the original context in which we were discussing basketball players, the sentence “I am tall” is false.

We can observe accommodation in the case of because-claims, as well. Consider (7), said out of the blue.

(7) Adam ate the apple because he was hungry.

Simply asserting the sentence makes it the case that the salient contrast is one between eating and other things Adam could have done with the apple, such as juggling it. That’s why it’s so natural to hear it as true. Imagine later saying (8).

interpretation of the first conjunct is different in C_1 and C_3 . All that follows from the acceptability of (2) is that the first conjunct is true with respect to both. This might be like evaluating “I am hungry” with respect to a context in which Sue is the speaker and is hungry, and then evaluating it with respect to a context in which Mary is the speaker and is hungry.

In that case, the skeptic is making the claim that (5a) is true with respect to a condition that (5b) does not satisfy—the one salient in C_1 —and also true with respect to a condition that (5b) does satisfy—the one salient in C_3 . But in that case, the skeptic still has not given us any reason for believing that in addition to the relation salient in C_3 , there is a further relation salient in C_1 that (5a) but not (5b) satisfies. So the skeptic’s case is no better if this option were actual. I’ll disregard it from now on.

(8) Adam ate the apple because Eve offered it to him.

Again, it seems natural to hear that sentence as saying something true. In order for that to be the case, the salient contrast must be one between Adam's eating the apple and someone (or something) else's eating it. So the mere assertions of (7) and (8) change the context so as to make each come out true. This, too, is an example of accommodation, and it shows that because-claims allow for accommodation in the normal way.

We now need to test for whether accommodation takes place in the interpretation of the conjunction (6). The crucial feature of accommodation is that it is not mandatory. It is usually open to the audience to resist the accommodation by explicitly rejecting the assertion to which the audience is asked to accommodate.

Imagine that, in response to (8), someone were to say (9).

(9) That's not true. After Eve gave him the apple, he might have juggled it.

We can hear (9) as true, which means that the proposition (8) expresses in this context is false. Accommodation has not taken place.

No such resistance is available to someone confronted with the conjunction at issue in this section.

(10) A: Why did the conductor warp?

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

A: # No it didn't. It only warped because a strong current went through it.

A's second utterance is unacceptable. So *B* must have spoken truly, even though *A* did her best to resist an accommodation. That means that the truth of the sentence *B* used does not depend on accommodation taking place. And therefore, the conjunction must be acceptable in a context because the truth-conditions of because-claims are not sensitive to a contextually variable relevance relation. If I can sustain this conclusion, I have established the second claim I urged against the skeptic. So let me consider what I take to be the most important challenge.

Irresistible Accommodation?

The argument crucially depends on the claim that accommodation can always be resisted. There are two ways to deny that. Suppose Sue and Mary are talking. Sue

says some things, including “I’m hungry.” Mary then begins speaking, and the first thing she says is “I’m not hungry.” Clearly, Mary’s utterance is about herself, not Sue. Hence, the context shifts to ensure that “I” is interpreted correctly. It is also obvious that Sue cannot resist this change in context by saying “No, I’m not”. Since Mary did not do anything before she started to speak, it’s also clear that her speaking forced the context-change. So it seems as if we have a case of accommodation (context-change as a result of an assertion) that cannot be resisted.

However, this is not an example of accommodation. Accommodation takes place only if the context-shift occurs in order to make the asserted sentence express a truth. That is obviously not what’s going on here. Any assertion of Mary’s would have made it the case that “I” is interpreted as referring to Mary, not Sue. Just imagine if Mary had said “I am a rhinoceros.” That, too, would have forced the shift in context, but the sentence would have been false.

A more promising example of irresistible accommodation is this. Suppose we’re in a department containing linguists and philosophers. It’s time for a department meeting, and all of the philosophers but none of the linguists have arrived. Consider the following dialogue.

- (11) A: Everybody is here.
B: The linguists aren’t here.
A: # No, everybody is here.

A’s second utterance is clearly unacceptable. This is a case of irresistible accommodation only if B’s utterance forces a context shift by accommodation. Whether this is the case depends on how exactly we theorize about the two claims A and B make. On one theory, B’s utterance forces a context-shift, but not by accommodation. On another, it forces a context shift by accommodation. Both views agree that the context changes so as to make A’s first utterance true, the second false. If that wasn’t the case, we wouldn’t have an example of context-shift at all. Thus, both views agree that when we evaluate A’s first utterance, the context is such that the quantifier ‘everybody’ is restricted to the philosophers. When we evaluate A’s second utterance, this is no longer possible. The theories disagree over how this happens.

According to the first theory, the domain of quantification with respect to which the quantifier is evaluated is constant throughout the exchange. But by one of several possible semantic or pragmatic processes, the quantifier is interpreted as saying the

same, in the context, as ‘everybody who is a philosopher’ would say in that context.²⁶

According to the second proposal, ‘everybody’ quantifies over everyone in the domain. The effect of restricted quantification is achieved by varying the domain with respect to which all utterances are evaluated. On this view, the domain with respect to which we evaluate *A*’s first utterance contains no linguists, so the utterance is true. The domain with respect to which we evaluate *A*’s second utterance does contain linguists, which is why the second utterance is false. This change in the context, moreover, is brought about by accommodation, because *B*’s utterance, when interpreted with respect to the first context would not be true. That’s because there are no linguists in the domain for “the linguists” to refer to. So to ensure that *B*’s utterance is true, the context must shift so that the domain contains linguists.

I think that we have good reason to reject this latter proposal.²⁷ Consider the sentence (12).

(12) In five departments, everybody is friendly with everybody else.

This sentence has a reading, indeed the most prominent one, according to which in each of five departments, the members of that department are friendly with one another, though they need not also be friendly with the members of the other departments. Thus, the sentence could truly describe a situation in which all of the biologists like each other, the philosophers like each other, and so on for the physicists, chemists, and linguists, though the biologists like neither the linguists nor anybody else outside their own departments, and likewise for the other departments.

The proposal on which the effect of restricted quantification is achieved by shrinking or expanding the domain cannot account for this reading. The initial phrase ‘in five departments’ quantifies over all five departments, so that the biologists, linguists, philosophers, chemists, and physicists are all in the domain. Hence, the phrase ‘everybody is friendly with everybody else’ is interpreted with respect to this large domain, and thus has to mean that any given member of one department likes every member of *every* department, not just their own. Hence, the shrinking and expanding domain

²⁶This could be a matter of a phonologically null but syntactically realized constituent of logical form, as proposed by Stanley and Szabo (2000). It could also be a matter of what Perry (1998) calls an unarticulated constituent, which is not syntactically realized but nonetheless influences what is said by the sentence. It could also be a matter of more pragmatically inflected processes, as proposed by, e.g., Bach (1994) and cois Recanatì (2002, 2004). The details do not matter here.

²⁷A similar argument can be found in Stanley and Szabo (2000, pp. 248-50). They use the argument to argue for a very specific view about how quantifiers work in context. However, I do not need to be committed to their particular approach in order to rule out the variable-domain option I’m concerned to argue against here. Any of the proposals I mentioned in note 26 would do for me.

cannot capture one of the readings of (12) and should therefore be rejected. For that reason, the initial example (11) is not a case of accommodation at all, and hence not irresistible accommodation, either. Skepticism fails.

Kitcher and Salmon's Critique SKEPTICISM

Before moving on to the second stage in my argument against the moderate, let me briefly pause to contrast my argument against SKEPTICISM with the famous critique in Kitcher and Salmon (1987).

They point out that SKEPTICISM is committed to the claim that almost any proposition could count as explanatory for some explanans in some context, and they say:

We take it that this result is counterintuitive. Indeed, we would view it as a *reductio* of van Fraassen's account of explanation.²⁸

However, this is just accusing the skeptic of her position.²⁹ And Kitcher and Salmon do not really treat it as a *reductio*, since they give arguments against this claim. They consider examples we do not wish to count as explanations, no matter what, such as (13).

(13) JFK died on 11/22/1963 because he was born under a Gemini Sun.

They claim that there is no context in which we would accept (13), even though its components are true, and they conclude that there must be invariant constraints on what can count as an explanation.

The skeptic can resist this argument. For one, it is unclear how compelling claims about which conversational situations *could not possibly* obtain are. It is surely a substantive claim—and if true an important observation—of contextualists in epistemology that the truth of our knowledge-ascriptions to a subject can vary, even if we hold the evidential facts fixed.³⁰ Before we were presented with the relevant cases, we were surely happy to grant that there could be cases where one ascription is more or less apt. But I think we would have resisted the stronger conclusion that there could actually be a difference in the truth-value of a knowledge ascription as we move from one context

²⁸Kitcher and Salmon (1987, p. 319).

²⁹Many other theorists who discuss van Fraassen's account of explanation level the same accusation. See, for example, Grimes (1987, p. 89), Hitchcock (1996, p. 399), and Sandborg (1998).

³⁰See DeRose (1992) for an early statement. An earlier, less appreciated, observation about the variable acceptability of knowledge ascriptions is in Austin's (1979).

to another. Claims about impossibility like the one Kitcher and Salmon are making should therefore be treated very cautiously.

Even if we grant them their claim, the skeptic can still resist the argument. Let me make the inference explicit. The premise is that (13) is infelicitous in any context we can imagine. Kitcher and Salmon conclude that it is ruled out by a context-invariant condition. This conclusion follows only if we accept the principle that if a sentence is unacceptable in every context, that must be due to specific semantic features of the sentence. And this principle is false. Consider the following parallel situation.

Simple theories that analyze causation as counterfactual dependence are committed to the claim that someone's birth caused their death. There is no context in which this is a natural thing to say, but that does not yet show that the theory is mistaken. The unacceptability of the causal claim may be due to requirements of informativeness, or perhaps constraints on specifically causal talk.

The skeptic cannot resist my argument by any of these methods. If successful, it establishes directly that the relevant semantic constraint obtains.

6 Constant Relevance Across Domains

I now turn to the second stage of my argument against the moderate, that the same relevance relation is used in interpreting because-claims from different domains. What I need to argue is that there are pairs of because-claims that come from different domains, but which address the same contrast.

With that in mind, consider (14) and (15).

- (14) a. Sue knocked on the door (rather than not) because she wanted to go inside.
b. Sue knocked on the door (rather than not) because it would have been rude not to.
- (15) a. After entering '2+2=,' the calculator showed '4' (rather than some other number) because it is wired in such-and-such way.
b. After entering '2+2=,' the calculator showed '4' (rather than some other number) because $2+2=4$.

I have made the contrast explicit in order to ensure that we are evaluating the because-claims in each pair with respect to the same contrast. I take it that all four of these

claims are true. Hence, at least in these examples, what it is to address a contrast cannot distinguish between the kinds of information offered in each pair.

One might worry at this stage that though I have made the contrast somewhat explicit, there are still subtle shifts in the contrast induced by the context. But this worry can be answered by the familiar anti-accommodation argument. If we conjoin the sentences in each of the pairs, the resulting conjunction is still true. And we have good reason for believing that accommodation does not take place, since we cannot resist the alleged accommodation by rejecting the second conjunct. So it's surely the case that we evaluate both because-claims in each pair with respect to the same contrast.

This argument does not show that moderation is false, given that we have left the position relatively vague. We have not made any assumptions about how many different ways of addressing a contrast the moderate contextualist wishes to countenance. What I have done is provide a strategy for testing specific claims the moderate might make. Let me state it generally.

Suppose she claims that there are different modes of addressing a contrast in two domains of inquiry D_1 and D_2 . We need to see if we can find a pair of because-claims that are interpreted with respect to the same contrast and that mention potentially explanatory information from D_1 and D_2 , respectively. If that is possible, we should reject the moderate's claim about D_1 and D_2 .

But even what I've said so far is sufficient to argue that causal theories of explanation are mistaken. Causal theories of explanation are committed to moderation. More specifically, a causal theory must be restricted to cases in which both the explanandum and the explanans can enter into causal relations. If a theory of explanation had to cover not just this case, but also cases in which the relationship between explanandum and explanans could not be causal, such as in mathematics, a causal theory of explanation would obviously fail.

I've argued here that this is precisely the case. A theory of explanation has to give a uniform account of what makes information explanatory across domains in which causation applies and domains in which it does not. Therefore, a causal theory of explanation must be mistaken.

7 Conclusion

I've argued that both skepticism and moderation fail. The conclusion of this argument is that there is a component to the truth-conditions of because-claims that is constant

across all of the fields of inquiry I've considered, and is quite possibly constant across all. Given this result, causal theories of explanation fail immediately.

But this does not answer all of our questions about explanation. For one, I have not said anything about what makes something a good explanation in a certain situation. It may well be the case that we distinguish between alternative, true because-claims as giving better or worse explanations. And one can have different attitudes towards such distinctions. One could maintain that there is a single account to be given here, as well, so that there is a domain-independent notion of being a good or better explanation.³¹ Alternatively, one could maintain that what counts as a good or better explanation is a far more situational matter. What counts as a good or better explanation might depend on the cognitive and technical resources available to us at a time. I take it this latter view is defended for example by Miller (1987), who holds that an explanation consists of adequately specified causal information, and who adds that what counts as adequate in a situation depends on the stage of inquiry we achieved in that situation. While I have maintained that an explanation cannot always consist in giving causal information, nothing I have said here takes a stand on whether the standards for good or better explanations are domain-transcendent, domain-specific, or perhaps even situation-specific.

Another question I have not spoken to is what I consider the foremost one right now. Straight-forward linguistic evidence suggests that our explanatory discourse involves questions besides why-questions: just as we can explain why something happened, so can we explain how it happened.³² Presumably, answering (at least some) why-questions and answering (at least some) how-questions have something important in common in virtue of which they are all explanations. Until we have that, we cannot be said to have an account of explanation. Having an account of what it takes to answer a why-question is an extremely important step along the way. But we should not confuse that with having reached our destination.

³¹Lewis (1986) tries to formulate this kind of position, arguing that all selection among bits of explanatory information is driven by completely general, broadly Gricean, principles.

³²Sylvain Bromberger first made this point in Bromberger (1992a).

References

- Achinstein, P. *The Nature of Explanation* (New York: Oxford UP, 1983).
- Austin, J. “Other Minds”. In *Philosophical Papers*, 76–116 (Oxford: Oxford UP, 1979).
- Bach, K. “Conversational Implicature”. *Mind and Language*, 9, (1994), 124–62.
- Boër, S. E. and W. G. Lycan. *Knowing Who* (Cambridge, MA: MIT Press, 1986).
- Braun, D. “Now You Know Who Hong Oak Yun Is”. *Philosophical Issues*, 16, (2006), 24–42.
- Bromberger, S. “An Approach to Explanation”. In *On What We Know We Don’t Know*, 18–51 (Chicago: The University of Chicago Press, 1992a).
- . “Why-Questions”. In *On What We Know We Don’t Know*, 75–100 (Chicago: The University of Chicago Press, 1992b).
- Cappelen, H. and E. Lepore. *Insensitive Semantics* (Malden, MA: Blackwell Publishing, 2005).
- Cartwright, N. *The Dappled World: A Study of the Boundaries of Science* (Cambridge, UK: Cambridge UP, 1999).
- Chomsky, N. *New Horizons in the Study of Language and Mind* (Cambridge, UK: Cambridge UP, 2000).
- DeRose, K. “Contextualism and Knowledge Attributions”. *Philosophy and Phenomenological Research*, 52(4), (1992), 913–929.
- Friedman, M. “Explanation and Scientific Understanding”. *The Journal of Philosophy*, 71(1), (1974), 5–19.
- Grimes, T. R. “Explanation and the Poverty of Pragmatics”. *Erkenntnis*, 27, (1987), 79–92.
- Groenendijk, J. and M. Stokhof. “Questions”. In J. van Benthem and A. ter Meulen, eds., *Handbook of Logic and Language*, 1055–1124 (Amsterdam: Elsevier and MIT Press, 1997).
- Hamblin, C. L. “Questions”. *Australasian Journal of Philosophy*, 36, (1958), 159–168.
- . “Questions in Montague English”. *Foundations of Language*, 10, (1973), 41–53.
- Harman, G. H. “The Inference to the Best Explanation”. *The Philosophical Review*, 74(1), (1965), 88–95.
- Hausman, D. M. *Causal Asymmetries* (Cambridge, UK: Cambridge UP, 1998).
- Hempel, C. G. “Aspects of Scientific Explanation”. In *Aspects of Scientific Explanation*, 331–496 (New York: The Free Press, 1965).

- Hitchcock, C. R. "The Role of Contrast in Causal and Explanatory Claims". *Erkenntnis*, 107, (1996), 395–419.
- Horgan, T. "From Supervenience to Superdupervenience: Meeting the Demands of a Material World". *Mind*, 102(408), (1993), 555–586.
- Kadmon, N. *Formal Pragmatics* (Malden, MA: Blackwell, 2001).
- 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).
- Kitcher, P. and W. Salmon. "Van Fraassen on Explanation". *The Journal of Philosophy*, 84(6), (1987), 315–330.
- 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).
- Miller, R. W. *Fact and Method* (Princeton: Princeton UP, 1987).
- Montague, R. "The Proper Treatment of Quantification in English". In P. Portner and B. H. Partee, eds., *Formal Semantics: The Essential Readings*, 17–34 (Malden, MA: Blackwell Publishing, 2002).
- Perry, J. "Indexicals, Contexts, and Unarticulated Constituents". In *Proceedings of the 1995 CSLI-Amsterdam Logic, Language and Computation Conference* (Stanford: CSLI Publications, 1998).
- Pietroski, P. M. "Meaning Before Truth". In G. Preyer and G. Peter, eds., *Contextualism in Philosophy*, 255–302 (Oxford: Oxford UP, 2005).
- Railton, P. "A Deductive-Nomological Model of Probabilistic Explanation". *Philosophy of Science*, 45, (1978), 206–226.
- . "Probability, Explanation, and Information". *Synthese*, 48, (1981), 233–256.
- Recanati, F. "Unarticulated Constituents". *Linguistics and Philosophy*, 25, (2002), 299–345.
- . *Literal Meaning* (Cambridge, UK: Cambridge UP, 2004).

- 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).
- Sandborg, D. “Mathematical Explanation and the Theory of Why-Questions”. *The British Journal for the Philosophy of Science*, 49(4), (1998), 603–624.
- Stanley, J. and Z. G. Szabo. “On Quantifier Domain Restriction”. *Mind and Language*, 15(2-3), (2000), 219–261.
- Steiner, M. “Mathematical Explanation”. *Philosophical Studies*, 34, (1978a), 135–151.
- Strevens, M. “The Causal and Unification Approaches to Explanation Unified—Causally”. *Noûs*, 38(1), (2004), 154–176.
- Van Fraassen, B. *The Scientific Image* (Oxford: Oxford UP, 1980).
- Woodward, J. *Making Things Happen* (Oxford: Oxford UP, 2003).