

ROUTLEDGE STUDIES IN CONTEMPORARY PHILOSOPHY

Contrastivism in Philosophy

Edited by
Martijn Blaauw



2 Causal Contextualism

Jonathan Schaffer

Causal statements are commonly made in some context, against a background which includes the assumption of some *causal field*. A causal statement will be the answer to a causal question, and the question ‘What caused this explosion?’ can be expanded into ‘What made the difference between those times, or those cases, within a certain range, in which no such explosion occurred, and this case in which an explosion did occur?’ Both causes and effects are seen as differences within a field. (Mackie 1974, 34–35)

Causal claims are context sensitive. For instance, if the engineer finds that the poor road conditions contributed to the accident, then it would be acceptable for her to say:

1. The poor road conditions caused the accident

Yet if the detective wants to focus on the drunk driver, then it would seem acceptable for him to deny 1 and instead say:

2. The poor road conditions didn’t cause the accident, it was the drunk driver

So much is commonplace. As Lewis notes:

We sometimes single out one among all the causes of some event and call it ‘the’ cause, as if there were no others. Or we single out a few as the ‘causes’, calling the rest mere ‘causal factors’ or ‘causal conditions’. . . We may select the abnormal or extraordinary causes, or those under human control, or those we deem good or bad, or just those we want to talk about. (1986, 162)

Yet, despite extensive studies of context sensitivity for other aspects of language such as knowledge ascriptions, there has been little discussion of the context sensitivity of causal claims. I will address three questions. In section 1, I will address the question of whether the context sensitivity of causal claims is partly semantic, or wholly pragmatic. I will argue—in a way familiar from arguments for epistemic contextualism—that the context

sensitivity of causal claims is partly semantic since it does not fully fit the pragmatic mold. In section 2, I will consider the question of whether causal claims are sensitive to contrasts, defaults, and/or models. I will argue that treating causal claims as sensitive to contrasts (for both cause and effect) does all the needed work. Finally in section 3, I will face the question—naturally arising from my answers to the first two questions—of how semantic sensitivity to contrasts might be implemented within an overall plausible semantic framework. This will turn out to be something of a puzzle. Accordingly, I must conclude that we do not yet have a clear understanding of context sensitivity as it arises for causal claims.

For those familiar with the discussion of the context sensitivity of knowledge ascriptions, it might be worth flagging two main respects in which the context sensitivity of causal claims will prove to differ. On the one hand, the intuitive data for context sensitivity is much stronger and more robust for causal claims, and includes specific phenomena that seem to have no counterpart with knowledge ascriptions (for instance, the matter of *selection* by which one causal factor is promoted to cause and the remainder demoted to background conditions). On the other hand, the semantic implementation of context sensitivity turns out to be far more problematic for causal claims, at least given the sort of contrastive views I advocate. This is because knowledge ascriptions only need a single source of contrasts which arguably can be read off the question under discussion. But causal claims need at least two separate sources of contrasts, and there is no obvious general procedure to recover the specific contrast applicable to the cause, or to recover the specific contrast applicable to the effect.

1 PARTLY SEMANTIC, OR WHOLLY PRAGMATIC?

Causal claims are context sensitive. That is, it may be acceptable for one speaker in one context to make a given causal claim, and acceptable for another speaker in another context to deny that very claim. This is uncontroversial. But what is controversial is whether such context sensitivity is a purely pragmatic affair, entirely explained by the extent to which the causal claim constitutes a cooperative contribution to the conversation at hand; or whether there is some semantic aspect to this context sensitivity. For instance, when the engineer finds that the poor road conditions contributed to the accident and says:

1. The poor road conditions caused the accident

And the detective denies 1 to lay the blame on the drunk driver, can it be that both the engineer and the detective still speak *truly*? Or must at least one of these characters (presumably the detective) be uttering a felicitous falsehood?

More precisely, what is at issue is the following thesis:

Causal Contextualism: A single causal claim can bear different truth values relative to different contexts, where this difference is traceable to the occurrence of ‘causes,’ and concerns a distinctively causal factor.

The first clause of *Causal Contextualism* characterizes the form of context sensitivity at issue: variation in truth value for a single sentence at multiple contexts of utterance. The remaining clauses try to ensure that this variation is arising for the right reason: not due to some other element of the sentence (perhaps all sentences contain other context sensitive elements), and not due to irrelevant features of ‘cause’ (such as tense and mood). This definition could perhaps use refinement, but should be clear enough to put to work.

1.1 The Context Sensitivity of Causal Claims

In order to evaluate *Causal Contextualism*, it will prove useful to provide a range of illustrations of the context sensitivity of causal claims. I do not claim that these illustrations exhaust all the context sensitivity of causal claims (I doubt they do), or that they must all receive a unified theoretical treatment (though I will offer one in terms of sensitivity to contrasts). Rather my purpose is to exhibit a family of striking and pervasive context sensitivities in causal discourse, in order to consider whether they fully fit the pragmatic mold. When I speak of “the context sensitivity of causal claims” in what follows, I should be understood as speaking of the sort (or sorts) of context sensitivity exhibited in these illustrations.

To begin, there is context sensitivity with respect to *causal selection*. It is part of causal discourse to promote some handful of factors to the status of cause, and to demote the remaining factors to the status of background condition.¹ This is the phenomenon seen in the case of the car accident above, and in Hart and Honoré’s example of the Indian famine:

The cause of a great famine in India may be identified by an Indian peasant as the drought, but the World Food Authority may identify the Indian government’s failure to build up food reserves as the cause and the drought as a mere condition. (1985, 35–36).

To provide one more illustration, the forest rangers would presumably promote the lightning strike to the status of cause for the forest fire, and would demote the presence of oxygen to the status of background condition. But now consider Putnam’s visiting Venusians: “Imagine that Venusians land on earth and observe a forest fire. One of them says, ‘I know what caused that—the atmosphere of the darned planet is saturated with oxygen’” (1982, 150).

So in particular, we can imagine a conversation among the Venusians in which the following claim was acceptable:

3. The presence of oxygen caused there to be a forest fire

Yet if we imagine a conversation among the forest rangers, 3 will surely be unacceptable in such a context. The forest rangers will deny that the presence of oxygen caused the fire. In this vein Hart and Honoré note:

In most cases where a fire has broken out the lawyer, the historian, and the plain man would refuse to say that the cause of the fire was the presence of oxygen, though no fire would have occurred without it: they would reserve the title of cause for something of the order of a short-circuit, the dropping of a lighted cigarette, or lightning. (1985, 11)

Further, there is context sensitivity with respect to *causal inquiry*. Causal claims are answers to ‘why’-questions, and differences in the preceding ‘why’-question may trigger differences in the acceptability of the resulting causal claims (van Fraassen 1980). For instance, if the question arises as to why John *kissed* Mary (perhaps we are wondering about John’s courage in matters of love), then a causal answer should explain why John offered Mary a kiss rather than, say, a hug or a handshake. On the other hand, if the question arises as to why John kissed *Mary* (perhaps we are wondering about John’s attraction to Mary), then a causal answer should explain why the person John kissed was Mary rather than, say, Suzy or Billy.

So in particular we can imagine a conversation in which we are presupposing that John loves Mary but questioning his romantic courage, in which the following claim might well be acceptable:

4. John’s boldness caused him to kiss Mary

Yet if we imagine a conversation in which we are presupposing John’s romantic courage, but questioning why he was attracted to Mary, an utterance of 4 may well be unacceptable. In such a context, one wants to hear about some feature of Mary (such as her sense of humor, or her flowing hair) that distinguishes her from her rivals.

Moreover, context sensitivity arises when there are *multiple alternatives* (Hitchcock 1996). For instance, suppose that the train switch has three settings. Setting it to *broken* will send the train down the broken track and on to disaster, setting it to *local* will send the train down the local track and slowly to the station, whereas setting it to *express* will send the train down the express track and swiftly to the station. The switch gets set to *local* and—just as expected—the passengers arrive slowly at the station. Did the switch’s getting set to *local* cause the passengers to arrive at the station (not to arrive slowly, just to arrive at all)? The answer seems to be: *it depends on which other option you had in mind*.

So in particular we can imagine a conversation in which the background assumption is that the switch was set to *broken*, and we are wondering why disaster was averted. In such a context the following claim should be acceptable:

5. The switch's getting set to *local* caused the passengers to arrive at the station

After all, if we were expecting the train to derail, learning that the switch got set to *local* should help us understand why things went otherwise. Yet if we imagine a context in which the background assumption is that the switch should be set to *express* (and leaving it on *broken* is not under consideration), then an utterance of 5 should be unacceptable. After all, in such a context we were already expecting the passengers to arrive at the station. The switch's getting set to *local* makes no difference.

1.2 Sentential Sensitivities

So far I have illustrated three sorts of context sensitivity for causal claims. It will also prove useful to display some other "nearby" sensitivities in causal discourse which are not sensitivities of a single sentence to context, but rather sensitivities between different sentences employing distinct but still coreferential event descriptions. The issue in these cases is how the differences in the event descriptions impact acceptability, and the hope is that these cases might shed light on how contextual differences impact acceptability.

With this in mind, consider the role that explicit 'rather than' clauses can play in causal discourse. For instance, in the train case above, one wants to say that the switch's being set to *local* rather than *broken* caused the passengers to arrive at the station:

6. The switch's getting set to *local* rather than *broken* caused the passengers to arrive at the station

But equally one wants to deny that the switch's being set to *local* rather than *express* caused the passengers to arrive at the station (they would arrive safely either way), by denying:

7. The switch's getting set to *local* rather than *express* caused the passengers to arrive at the station

Yet unless one has an implausibly fine conception of events, it seems that the switch's getting set to *local* rather than *broken*, and the switch's getting set to *local* rather than *express*, pick out the same event under a different description. It is not as if the switch got set twice.

Moreover, it is not as if the switch's being set to *local* rather than *express* made no difference. Its being set to *local* caused the passengers to arrive at the station slowly rather than swiftly:

8. The switch's getting set to *local* rather than *express* caused the passengers to arrive at the station slowly rather than swiftly

And again, unless one has an implausibly fine conception of events, it seems that the passenger's arriving at the station, and the passenger's arriving at the station slowly rather than swiftly, pick out the same event. How are the 'rather than' clauses impacting acceptability, if the same events are picked out either way?

Or consider the role that more specific descriptions of events can play in causal discourse. To borrow an example from McDermott (1995), one might deny that McEnroe's tension caused him to serve, but accept that it caused him to serve awkwardly:

- 9. McEnroe's tension caused him to serve
- 10. McEnroe's tension caused him to serve awkwardly

That is, while 9 seems unacceptable, 10 is fine. One wants to say: the tension didn't matter to *whether* he served but only to *how*. Yet unless one has an implausibly fine conception of events, it seems that McEnroe's serving *just was* his serving awkwardly. We are just discussing a single serve.²

Indeed, as the following example from Achinstein (1975) shows, merely shifting the locus of focus within the event nominal can control acceptability:

- 11. Socrates's DRINKING HEMLOCK at dusk caused his death
- 12. Socrates's drinking hemlock AT DUSK caused his death

11 seems acceptable, but 12 does not. One wants to say: what Socrates drank mattered, when he drank it did not. Again, unless one has so fine a conception of events that focal differences can make for event differences, the same event of Socrates's drinking hemlock at dusk is described in both 11 and 12, merely with a difference in emphasis. There was just one drinking.

1.3 The Invariantist Orthodoxy

The question is whether the contextual and sentential sensitivities just illustrated indicate any sort of *semantic* context sensitivity in causal discourse (as per *Causal Contextualism*), or whether they can be fully explained via conversational pragmatics. Perhaps all one sees is differences in the extent to which a given causal claim is a cooperative contribution to different conversations.

To the extent that there is an orthodox view in the current literature, it is the view that the context sensitivity of causal claims is entirely pragmatic. This view denies *Causal Contextualism* without denying the sensitivity "data," instead positing a purely pragmatic explanation for this data:

Causal Invariantism: It is not the case that a single causal claim can bear different truth values relative to different contexts, where this

difference is traceable to the occurrence of ‘causes,’ and concerns a distinctively causal factor. Causal claims are context sensitive in their acceptability, but the context sensitivity of causal claims is a wholly pragmatic phenomenon.

The first sentence of *Causal Invariantism* is the denial of *Causal Contextualism*, and the second sentence adds the acceptance of context sensitivity plus the posit of a purely pragmatic explanation for such context sensitivity.³

For the invariantist, causal discourse involves a preselective semantics for some egalitarian and unselective notion of *being a causal factor*. In the forest fire case above, both the lightning strike and the presence of oxygen should equally qualify. Indeed presumably all the positive causal claims made in sections 1.1–1.2, even if unacceptable in the context at hand, will count as true. In this vein, Lewis—while defending a counterfactual analysis of causation—clarifies that he is “concerned with the prior question of what it is to be one of the causes (unselectively speaking). My analysis is meant to capture a broad and nondiscriminatory concept of causation” (1986, 162).⁴

The invariantist then layers a selective pragmatics for *being a salient causal factor* (sometimes expressed as *being “the cause”*) over her preselective semantics. In the forest fire case, our interests and background expectations will determine which of the many “causes” gets selected as salient. Along these lines, Mackie speaks of causal selection as “reflecting not the meaning of causal statements, but rather their conversational point” (1974, 35), and Lewis explicitly associates causal selection with Gricean conversational pragmatics:

There are ever so many reasons why it may be inappropriate to say something true. It might be irrelevant to the conversation, it might convey a false hint, it might be known already to all concerned, and so on (Grice 1975). (2004, 101; cf. Bennett 1995, 133)

Of course it is uncontroversial that there are pragmatic phenomena in discourse, and *a fortiori* uncontroversial that there are pragmatic phenomena in causal discourse. The question is whether pragmatics can fully explain the contextual and sentential sensitivities exhibited. To my knowledge no invariantist has ever tried to spell out the pragmatic explanations in any detail, or do much more than allude to the prospect of some Gricean story.

From a wider perspective, *Causal Contextualism* might be counted as orthodoxy, and rooted in Mill’s groundbreaking discussion of causal selection. For Mill is a *revisionist* about causal discourse. He thinks that our causal claims are shot through with selection effects. He merely regrets this as unscientific and deserving of excision:

Nothing can better show the absence of any scientific ground for the distinction between the cause of a phenomena and its conditions, than

the capricious manner in which we select from among the conditions that which we choose to denominate the cause. (1950, 244)

So when it comes to a purely descriptive account of our causal concept, Mill looks to be on the contextualist side. Contextualism also has roots in Hart and Honoré's discussion of the role of causation in the law: "The contrast of cause with mere conditions is an inseparable feature of all causal thinking, and constitutes as much the meaning of causal expressions as the implicit reference to generalizations does" (1985, 12). Contextualism has further roots in van Fraassen's (1980) discussion of the context sensitivity of explanation, insofar as both causal and explanatory claims are understood as triggered by contrastive why questions. And contextualism seems to have attracted a fairly wide range of contemporary theorists, including Hitchcock (1996), Woodward (2003), Maslen (2004), Menzies (2004 and 2007), Schaffer (2005a and 2010), Hall (2007), and Northcott (2008). So perhaps a new orthodoxy is (re-)forming.

1.4 Against Invariantism

Evidently there are pragmatic phenomena in discourse, and *a fortiori* there are pragmatic phenomena in causal discourse. The question is whether pragmatics serves to fully explain the context sensitivity of causal claims. I will now offer three connected arguments that conversational pragmatics cannot fully explain this context sensitivity. (Given that *Causal Invariantism* is the main alternative to *Causal Contextualism*, these arguments are indirectly arguments for the contextualist alternative.)

The first argument is that *no known pragmatic mechanism handles the cases*. So suppose that there is a lightning strike and a forest fire breaks out, and that a forest ranger utters 3, citing the presence of oxygen as causing the fire. Given that this is unacceptable, and given that the pragmatic explanations available are going to involve something like Gricean maxims—and in particular floutings of Gricean maxims which produce false implicatures—one can ask which Gricean maxim is flouted.⁵ The only Gricean maxim which seems applicable is Relevance. Indeed it seems clear that the remaining Gricean maxims—namely Quality, Quantity, and Manner—need not be flouted at all. After all, the forest ranger might have excellent evidence for there being oxygen present and for it being a factor, she has been informative, and she has not spoken in an overly prolix or otherwise marked manner.

So I take it that the only known pragmatic mechanism that might be operative in this case is Relevance, and that the invariantist will say that when the forest ranger is inquiring into the causes of the forest fire, she is *presupposing* that oxygen is present, and wondering about *what ignited the oxygen*. So citing the presence of oxygen is failing to speak to the question under discussion, and hence flouts Relevance.⁶

But this pragmatic explanation—which looks like the only one available in anything like a Gricean framework—is inadequate, because floutings of Relevance produce a distinctive feel not found in the illustrations, which is a not a feeling of falsity but merely of irrelevance. Thus consider Kierkegaard's (1978, 50) parable of the madman who repeats “Bang! The earth is round!” at every turn. One is inclined to label the man mad and his utterances irrelevant, *but there is no feeling that the man has said anything false*. We clearly recognize what the madman keeps repeating as an irrelevant *truth*. (We also recognize that prefacing every utterance with “Bang!” is a bit odd, but leave that aside.) No one, on considering Kierkegaard's madman, should feel any inclination to reject the claim that the earth is round.

Matters may be clearest in the examples with sentential differences from section 1.2. For instance, on the pragmatic view 7 and 9 are literally true but merely irrelevant:

7. The switch's getting set to *local* rather than *express* caused the passengers to arrive at the station
9. McEnroe's tension caused him to serve

Yet neither seems like an irrelevant truth which simply did not bear mentioning. Instead both feel more like falsehoods. The switch's getting set to *local* rather than *express* did not cause the passengers to arrive at the station—it made no difference whatsoever to whether the passengers arrived at the station, because both settings are stipulated to result in this same outcome. Likewise McEnroe's tension did not cause him to serve—it made no difference whatsoever as to whether he served, because he was set to serve anyway. The distinctive feel of irrelevant truth is absent.

The second argument against pragmatic explanations is that *speakers assert the negations*. Ordinary speakers will not only refuse to assert claims like 3 (in the context of the forest rangers); they will go so far as to assert its denial:

13. The presence of oxygen did not cause there to be a forest fire, what caused the fire was the lightning

(The more sophisticated speaker may then clarify that the presence of the oxygen was a “mere background condition” or something of that ilk.) Likewise in the case of the car accident, the detective who wants to focus on the drunk driver will *deny* that the poor road conditions caused the accident, as per 2. Floutings of Relevance will at best explain a refusal to assert 1 or 3. They will not explain a willingness to assert the negation as seen in 2 or 11, since (by the lights of invariantists) the negations are equally irrelevant and false to boot!⁷

Indeed the first two arguments against pragmatic explanations connect. When the pragmatic explanation involves a mere flouting of Relevance, the

assertion will have the feel of an irrelevant truth. This is why there will be no temptation to assert the negation (an irrelevant falsehood). Consider again Kierkegaard's madman. No sane and minimally informed speaker would go so far as to assert the negation: "The earth is not round."

Third and finally, *cancellation does not help*. The main test for conversational implicatures is that they (unlike semantic entailments) are cancelable. For instance, if I say of a job candidate that she has excellent handwriting, I can block the implicature that she is a poor philosopher by saying "but I don't mean to suggest that she is a poor philosopher" (I may then go on to discuss her philosophical genius). None of the causal cases pass the test. Thus consider, in the context of the forest rangers:

14. The presence of oxygen caused there to be a forest fire, but I don't mean to suggest that the lightning strike played no role

This hardly seems any more acceptable, despite the cancelation of any potential implicature that the lightning strike played no role. Or consider:

15. McEnroe's tension caused him to serve, although I don't mean to suggest that it was the only factor involved

These attempts at cancelation hardly seem to salvage acceptability. With 15, one wants to say that regardless of which factors did cause McEnroe to serve, his tension was not among them.

These three arguments thus constitute a *prima facie* case against *Causal Invariantism* and thereby a *prima facie* case for its main competitor, *Causal Contextualism*. I do not mean to suggest that these arguments are decisive. The invariantist can always challenge the "data" or try to introduce new pragmatic mechanisms to better explain it, and there are also alternatives to consider between *Causal Invariantism* and *Causal Contextualism* (for instance, perhaps some of these judgments should be explained away as performance errors of some sort). But I do mean to suggest that the invariantist orthodoxy, which assimilates the context sensitivity of causal claims to Gricean implicatures, is implausible. As soon as one tries to spell out the details of the Gricean story, it emerges that the context sensitivity of causal claims does not fit the pragmatic mold. Perhaps Mill was right from the start in his descriptive claim about our causal concept.

2 WHAT SHIFTS?

So far I have argued that the context sensitivity of causal claims is partly semantic, or at the least is not wholly a matter of Gricean conversational pragmatics. But leaving this aside, there is a largely independent question of which contextual parameters causal claims are sensitive to. What shifts

with context? That is, what gears of the contextual machinery are engaged by the illustrative cases above, on which the acceptability of the relevant causal claim turns?

To illustrate the sort of question I am asking, consider a simple indexical like 'I.' No serious account of 'I' could rest with the claim that it exhibited context sensitivity, or even with the claim that it exhibited semantic context sensitivity. There is the further question of which contextual parameters 'I' is sensitive to. In this case the answer is straightforward: the semantic value (content) of 'I' is sensitive to the contextual factor of *who is speaking*. (The reason why it is acceptable for me but not you to say 'I am Jonathan Schaffer' is that when I say it the 'I' refers to me and so the claim is true, but when you say it the 'I' refers to you and so the claim is false.) Presumably the context sensitivity of 'cause' is a sensitivity to some other factor or factors, but which?

Such a question is largely independent of the previous question as to whether the sensitivity is pragmatic or semantic, but not completely independent. For if 'cause' is sensitive to a given factor, then there must be a parameter at the relevant pragmatic or semantic level tracking this factor. With 'I,' given that it is sensitive to who is speaking, and given that this is a semantic matter, then there must be a semantic level parameter tracking who is speaking. With 'cause,' I will be arguing that it is sensitive to contrast, and given that this is a semantic matter (as previously argued), then there must be a semantic level parameter tracking the contrasts. I will return to this matter in section 3.

2.1 Contrasts, Defaults, and Models

What are causal claims sensitive to? It turns out that there are at least three different—albeit compatible and not wholly distinct—sorts of answers that one finds in the literature. One answer, which I will be defending (and which is defended in various forms in Hitchcock 1996; Woodward 2003; Maslen 2004; Schaffer 2005a; and Northcott 2008), is that causal claims are sensitive to *contrasts*. What shifts with context are the contrasts in play, where contrasts are specific possible alternatives to actual events. Actually there are at least four versions of contrastivism that are found in the literature, concerning whether one is looking at a contrast for the cause (c^*), for the effect (e^*), for both, or for each event in the set of events under consideration (V^*):

Cause-Contrast: c rather than c^* causes e

Effect-Contrast: c causes e rather than e^*

Double-Contrasts: c rather than c^* causes e rather than e^*

Total-Contrasts: c causes e relative to V^*

(One might also work with a set C^* of contrasts for the cause and/or a set E^* of contrasts for the effect, but I will suppress this complication for simplicity.)

The contrastive view can—though it need not—be plugged into a simple counterfactual test for causation by replacing the supposition of the

nonoccurrence of c or e (or of any intermediaries or other events involved in the account), with the supposition of the occurrence of the associated contrast. So for instance—at least as a decent gloss of *Double-Contrasts*—one might hold that c rather than c^* causes e rather than e^* iff (roughly) if c^* had occurred then e^* would have occurred. I will be defending *Double-Contrasts* (though I would be equally happy with *Total-Contrasts*—what is crucial is just that we have contrasts for both cause and effect; further contrasts might also prove useful). The counterfactual test just offered will prove useful insofar as it—together with certain assumptions about which contrasts are relevant in which contexts—will allow one to use *Double-Contrasts* to test truth values for causal claims.

But a different answer (supported by Menzies 2004; Hitchcock 2007; and Hall 2007) is that causal claims are sensitive to *defaults*. What shifts with context are which outcomes count as the “normal” or “default” behavior of the system under consideration, and which count as “abnormal” or “deviant” behavior. It is theoretically possible to assign defaults to a range of possible outcomes for the cause, for the effect, for both, or for every event under consideration (just as with contrasts), but all the defaultists in the literature have worked with the idea that defaults are assigned for all events under consideration:

Default: c causes e relative to *Def*.

Def is a function from each event under consideration to a range of “default” outcomes associated with that event (the actual event might be a default outcome or a deviant outcome).

A guiding idea behind *Default* is that causes and effects are conceptualized as deviations from some sort of natural state (Maudlin 2004). This idea can, for instance, be plugged into a simple counterfactual test by treating the nonoccurrence suppositions as reversions to default behavior. So for instance, where *Def* assigns a single default outcome for both c and e , one might hold that c causes e iff if *Def*(c) had occurred then *Def*(e) would have occurred.⁸

And yet a third answer (found in Menzies 2004; Halpern and Pearl 2005; Hitchcock 2007) is that causation is relative to an eligible causal model of the situation:

Model: c causes e relative to *Mod*.

Mod may be a set of variables and structural equations as in Pearl (2000), or a set of objects assumed to form a closed system plus a set of governing laws as in Menzies (2004). This is the natural reading of any theorist in the causal modeling tradition who gives an account of when one variable in a model is causally related to another variable in a model, while allowing (as is usually allowed) that there are worldly situations for which there are multiple eligible causal models with diverging causal verdicts.

The various contrastive proposals, *Default*, and *Model* are not wholly distinct, and indeed—at least on their leading implementations—can be ordered in strength as follows:

<i>Cause-Contrast</i> <i>Effect-Contrast</i>	<i>Double-Contrasts</i>	<i>Total-Contrasts</i>	<i>Model</i>	<i>Default</i>
-------------------------------------------------	-------------------------	------------------------	--------------	----------------

Working backwards, *Default*—at least as implemented in Hall (2007) and Hitchcock (2007)—is a strict addition to *Model*, since default structure is given by adding the *Def* function to Pearl models, augmenting the variables and structural equations of Pearl models with a function from each variable to some subset of its allotted values that are to count as its default settings. So implemented, default relativity might be understood as relativity to *augmented models* with an added *Def* function. And *Model* is an addition to *Total-Contrasts*, insofar as models include variables with a range of allotted values, which range is a contrast space for the event modeled by the variable.⁹ *Model* adds a further relativity to other aspects of causal models beyond the range of allotted values for variables, namely the choice of events modeled by variables and the structural equations over the variables. *Total-Contrasts* adds to *Double-Contrasts* a further relativity to contrasts for events under consideration other than cause or effect, and *Double-Contrasts* adds to both *Cause-Contrast* and *Effect-Contrast* a relativity to contrasts for the other side of the causal relation.

This means that *Total-Contrasts* can be thought of as partial model relativity. *Total-Contrasts* can be thought of as relativity to the range of allotted values for the variables, *without* relativity to the remaining aspects of the model, namely which events are represented by variables, and what structural equations hold over these variables. With respect to the events represented, it is natural to think that there is an objective fact as to which events are out in the world to be represented. Models which—for the sake of tractability—do not represent all the events idealize at their peril. With respect to the structural equations, these are generally supposed to hold objectively, representing the counterfactual facts as to what would lead to what. Fix which variables are modeled and what range of values they are allotted, and there is a right choice of structural equations. Any relativity to “a different choice” of structural equations is at best mistaken. Perhaps *Total-Contrasts* thus captures the element of truth in *Model*, while avoiding the other implausible aspects of model relativity. Though I will largely work with *Double-Contrasts* in what follows, my openness to the strengthened thesis of *Total-Contrasts* largely stems from this connection to causal models.

Note also that contrastivity, default-relativity, and model-relativity are compatible, and so one might endorse any combination thereof, including:

Contrast-Default-Model: c rather than c^* causes e rather than e^* relative to *Def* and *Mod*

Though—at least on the leading implementations of these ideas, on which they are ordered in strength as per above—such combinations are not genuinely new options. Given that default relativity includes model relativity (and thereby includes contrast relativity), *Contrast-Default-Model* is just *Default* by another name.

Note further that these options are hardly exhaustive. They are merely the main options that have been considered in the literature. One could also endorse an ambiguity thesis on which ‘cause’ can express a plurality of these candidates (Hitchcock 2003). That said, I will be arguing that contrastivity—and specifically *Double-Contrasts*—suffices to explain the context sensitivity of causal claims, so there seems no need (at least with respect to the cases currently under discussion) for anything further or stronger, or for any posited ambiguity. (Though again I am officially neutral between *Double-Contrasts* and *Total-Contrasts*.)¹⁰

I will further argue (section 2.3) that contrast sensitivity is specially rooted in the theoretical roles that causation plays. And I will argue (section 3) that there is independent linguistic reason to think that contrasts are elements of conversational context, and so are available as a contextual parameter to connect with causal claims. This claim does not carry over to defaults or models (or other arbitrary proposals). So the contrastive view also seems uniquely well situated with respect to linguistic implementation. Alas, I will also be arguing that the contrastive view faces linguistic difficulties as well, so I must be wary of claiming any ultimate advantage on this last matter.

2.2 Context Sensitivity as Contrast Sensitivity

I will now argue that contrastivity—and specifically *Double-Contrasts* on which causation is a relation of the form c rather than c^* causes e rather than e^* —serves to explain the context sensitivity of causation (section 1.1), and the nearby sentential sensitivity (section 1.2). Recall that the context sensitivity of causation, at least in the form under discussion, encompasses:

- causal selection (as illustrated by whether or not the presence of oxygen is said to cause there to be a forest fire)
- causal inquiry (as illustrated by the different causal answers appropriate for the questions of why *John* kissed *Mary*, why John *kissed* *Mary*, and why John kissed *Mary*)
- multiple alternatives (as illustrated by the train switch with the *broken*, *local*, and *express* settings)

And the nearby sentential sensitivity, at least in the form under discussion, encompasses:

- ‘rather than’-clauses (as illustrated by ‘the train switch being set to *local* rather than *express*’ as opposed to ‘the train switch being set to *local* rather than *broken*’)

- event specifications (as illustrated by ‘McEnroe’s serving’ versus ‘McEnroe’s serving awkwardly’)
- focus shifts (as illustrated by comparing ‘Socrates DRINKING HEMLOCK at dusk’ with ‘Socrates’s drinking hemlock AT DUSK’)

I will now argue that these sensitivities are all connected manifestations of an underlying contrast sensitivity in causal discourse.¹¹

It might help to start with the focus shift cases, since these are perhaps clearest in terms of the theoretical treatment required. Focus (at least of the sort exhibited in the cases at hand) is *contrastive focus*. Returning to 11 and 12:

11. Socrates’s DRINKING HEMLOCK at dusk caused his death
12. Socrates’s drinking hemlock AT DUSK caused his death

In 11, ‘Socrates’s DRINKING HEMLOCK at dusk’ is naturally interpreted as *c*: Socrates’s drinking hemlock at dusk, rather than *c**: Socrates’s drinking wine at dusk (or some other salient alternative to drinking hemlock); whereas in 12, ‘Socrates’s drinking hemlock AT DUSK’ is naturally interpreted as *c*: Socrates’s drinking hemlock at dusk, rather than *c**: Socrates’s drinking hemlock at dawn (or some other salient alternative to occurring at dusk). Indeed such a contrastive treatment falls out of leading linguistic treatments of focus such as Rooth’s (1992) alternative semantics, on which the focus semantic value of an expression is the result of replacing the focused constituent with the set of contextually salient options.¹² The difference between 11 and 12 is *not* between the actual events denoted, but between the contrasts selected.

Strictly speaking 11 and 12 only call for contrasts to the cause, as per *Cause-Contrast*. But it is easy to see that the same pattern can be found on the “effect side” as well, as seen in:

16. Socrates’s DRINKING HEMLOCK at dusk at dusk caused HIS DEATH at dawn
17. Socrates’s DRINKING HEMLOCK at dusk caused his death AT DAWN

16 seems acceptable but 17 does not, and these differ only in the locus of focus on the effect side. One wants to say that Socrates’s drinking hemlock rather than wine (as per the contrastive interpretation on the cause side) made a difference to whether or not he died, but not to when he died—had Socrates drunk wine he would have survived through the relevant time. (Though if the context is an unusual one in which the alternative of Socrates dying at a ripe old age is salient, then 17 should become acceptable. This is further confirmation of the way in which the contextual salience of contrasts controls acceptability.)

The shifts in ‘rather than’-clauses—which are just overt contrastives—clearly follow the same pattern. Indeed 11, 12, 16, and 17 can all be

rephrased—preserving the patterns of acceptability—with ‘rather than’-clauses concerning the focused alternatives instead of focus. Or to return to the train cases, recall 8 (which has ‘rather than’-clauses for both cause and effect):

8. The switch’s getting set to *local* rather than *express* caused the passengers to arrive at the station slowly rather than swiftly

This is acceptable since the difference between *local* and *express* is what made the difference between a slow and swift arrival. But vary either of the ‘rather than’-clauses to lose difference making and the result is unacceptable:

18. The switch’s getting set to *local* rather than *broken* caused the passengers to arrive at the station slowly rather than swiftly
 19. The switch’s getting set to *local* rather than *express* caused the passengers to arrive at the station slowly rather than suffer a derailling

After all, with 18 the passengers were not going to arrive swiftly whether the switch was set to *local* or *broken* (it is not as if setting the switch to *broken* would have sped up their arrival!) And with 19 the passengers were not going to suffer a derailling whether the switch was set to *local* or *express* (either way they are safe). Moreover, the ‘rather than’-clauses can still be re-correlated to re-gain difference making, with acceptability regained, as in:

20. The switch’s getting set to *local* rather than *broken* caused the passengers to arrive at the station slowly rather than suffer a derailling

Given that the ‘rather than’-clauses are overtly specifying the relevant contrasts (either directly providing the value of c^* and e^* , or—perhaps better—describing c and e in ways that naturally generate values for c^* and e^*), this is further direct evidence for an underlying contrast sensitivity in causal discourse.

Shifting to the event specificational differences, the very same pattern emerges. In the case of McEnroe’s serve, the underlying contrastive causal truths (made explicit via overt ‘rather than’-clauses) is as follows:

21. McEnroe’s being tense rather than calm caused his serving awkwardly rather than gracefully

While the following is a contrastive causal falsehood:

22. McEnroe’s being tense rather than calm caused his serving rather than standing still

All that needs to be added is that describing the effect event as a “serving”—as in 9—invites a contrast such as a standing still, and so invites

an interpretation via the falsehood of 22. But describing the effect event more specifically as a “serving awkwardly”—as in 10—invites a contrast such as a serving gracefully, and so invites an interpretation via the truth of 21. Thus the difference in acceptability between 9 and 10 is naturally explained on a contrastive treatment.

So far I have argued that the sentential sensitivity of causal claims (section 1.2) is due to an underlying contrast sensitivity. It remains to show that the contextual sensitivity of causal claims (section 1.1) evinces the same underlying pattern. It might help to start on this point with the multiple alternatives seen in the train case, by reconsidering:

5. The switch’s getting set to *local* caused the passengers to arrive at the station

The “data” observed in section 1.1 was that the acceptability of 5 seemed to vary with which alternative (*express* or *broken*) was salient. Assuming that the contextual alternative to arriving at the station is derailling, then on the contrastive treatment 5 is equivalent to the false contrastive claim 19 in contexts in which *express* is the salient contrast to the cause, and equivalent to the true contrastive claim 20 in contexts in which *broken* is the salient contrast.

Or consider sensitivity to the causal inquiry, as it impacts the acceptability of:

4. John’s boldness caused him to kiss Mary

Note that different causal inquiries are associated with different contrastive why-questions. These generate contrasts on *e*, and generate different slates of permitted answers, which generate contrasts on *c*. Think of the contrasts as the contextually permitted answers to the twofold question ‘What happened, and why?’, where the ‘What happened?’ provides the space of salient options for the effect and the ‘why?’ provides the space of salient options for the cause.

So if we are questioning John’s romantic courage the ‘What happened?’ aspect of the causal inquiry might present the options of (i) John kissed Mary, and (ii) John merely waved goodnight, and the ‘why?’ aspect of the causal inquiry might present the options of (iii) John is bold, and (iv) John is timid. In such a context 4 will be equivalent to the following contrastive truth:

23. John’s being bold rather than timid caused him to kiss Mary rather than merely waving goodnight

Yet if we are not questioning John’s romantic courage but instead questioning why he chose to kiss Mary, the ‘what happened?’ aspect of the causal

inquiry might present the options of (i) John kissed Mary, and (ii*) John kissed Billy. In such a context 4 will be equivalent to the following contrastive falsehood:

24. John's being bold rather than timid caused him to kiss Mary rather than Billy

Or at least, 24 is false given that John's preference for Mary over Billy is not a matter of boldness. (If John is a confirmed homosexual who is boldly experimenting with his sexuality, then 4 should become acceptable. This is further confirmation of the way in which the contextual salience of contrasts controls acceptability.)

Finally, returning to causal selection, recall how this impacts the acceptability of claims such as:

3. The presence of oxygen caused there to be a forest fire

The question is *why* the Venusians naturally promote the presence of oxygen to the status of a cause while the forest rangers naturally demote the presence of oxygen to the status of a mere background condition. A natural first thought is that, for the Venusians, there is a salient alternative to the presence of oxygen: *the absence of oxygen*. But for the forest rangers no alternative to the presence of oxygen is salient. For the forest rangers the presence of oxygen is simply presupposed.

So understood, there is a single contrastive truth in play:

25. The presence rather than absence of oxygen caused there to be a forest fire rather than no fire

The reason why 3 is acceptable in the context of the Venusians is because it is equivalent to 25, since the absence of oxygen is a relevant alternative for them. But since the forest rangers recognize no salient alternative to the presence of oxygen, 3 does not receive any such interpretation (nor is it obvious what if any interpretation it should receive).¹³

On this treatment, causal selection stems from different background expectations which generate different causal inquiries. The forest rangers are presupposing that oxygen is present, and in effect asking what ignited the oxygen. The information about the presence of oxygen does not answer their question. The Venusians on the other hand are presupposing that lightning strikes are present, and in effect asking what the lightning strikes ignited. The information about the presence of oxygen answers their question. Overall what seems to be governing selection is the causal inquiry and its attendant possible answers (the contrasts). In causal selection what varies in a "capricious manner" (as Mill says) is which contrasts are in play

in a given context, but what is predictable is what counts as the cause given the contrasts.

The above account of causal selection is essentially Mackie's view, on which a "causal statement will be the answer to a causal question" (1974, 34), and on which "causes and effects are seen as differences within a field" (1974, 35). The elements of the field are the contextually determined background conditions. Indeed, as I have suggested elsewhere (Schaffer 2005a), Mackie's view is the only plausible account of selection in the literature. If so then selection requires use of contrasts.¹⁴

Putting all of this together, *Double-Contrasts* seems capable of explaining all the contextual and sentential sensitivities of causation under discussion, and doing so in a unified and elegant way. Thus I would conclude that what shifts with context are the salient contrasts to the cause and to the effect.

2.3 Theoretical Motivations for Contrast Sensitivity

I have just argued—as an inference to the best explanation for the context sensitivity of causal claims—that causation is a contrastive relation, of the form c rather than c^* causes e rather than e^* . This conclusion may be buttressed by considering the theoretical roles of causation as a relation of *difference making*, as connected to *agential manipulation*, and as supporting *explanation*. Any relation that plays these roles needs contrastive structure.

As to difference making, recall what Lewis says in connecting counterfactuals to causal reasoning: "We think of a cause as something that makes a difference" (1986, 160–161). Lewis goes on to think of a cause as something whose occurrence or nonoccurrence makes a difference to the occurrence or nonoccurrence of the effect. But seen this way it should be evident that the notion of difference making is a contrastive notion. The contrasts articulate what the salient differences are. And it should be evident that the Lewisian notion of difference making involving nonoccurrence is just one of many ways of making a difference. There can also be a difference between a cause and an alternative to it (other than nonoccurrence), with respect to the effect versus an alternative to it (other than nonoccurrence).

Moreover, there is reason to think that the very idea of a "nonoccurrence" which Lewis appeals to is itself implicitly contrastive, in the sense that nonoccurrence suppositions take us *to the contextually salient alternative* (Schaffer 2005a). In this vein consider 'If John had not kissed Mary. . .'—one naturally imagines someone else doing the kissing. But consider 'If John had not *kissed* Mary. . .'—one naturally imagines something like a chaste handshake; or instead consider 'If John had not kissed *Mary*. . .'—now one naturally imagines John kissing someone else. Thus

the expressions ‘if c had not occurred’ and also ‘then e would not have occurred’ in Lewis’s counterfactual account are naturally read as equivalent to ‘if c^* had occurred’ and ‘then e^* would have occurred,’ where c^* and e^* are the contextually salient contrasts to c and e respectively.

As to agential manipulation, everyone accepts that there are connections between causation and the notions of intervention, manipulation, and agency. Never mind in which directions the connections run—all that matters here is that these notions are interconnected. Now the notion of manipulation seems patently contrastive, as Woodward explains:

Any manipulation of a cause will involve a change from one state to some specific alternative, and how, if at all, a putative effect is changed under this manipulation will depend on the alternative state to which the cause is changed. Thus, if causal claims are to convey information about what will happen under hypothetical manipulations, they must convey the information that one or more specific changes in the cause will change the effect (or the probability of the effect). This in turn means that all causal claims must be interpretable as having a contrastive structure. (Woodward 2003, 146)

So it seems that causation must embody some sort of sensitivity to alternative courses of action (“hypothetical manipulations”) if it is to properly connect to agency.

Finally, causation is widely thought to back explanation, and explanation has itself been argued to be contrastive (van Fraassen 1980; Garfinkel 1981). For instance, the explanation for why John kissed Mary rather than merely waving goodnight to her might differ from the explanation for why John kissed Mary rather than Billy. Or to express the matter with focus: the explanation for why John *kissed* Mary might differ from the explanation for why John kissed *Mary*. Given that causation serves to back explanation, it is most natural to posit that causal relations have the same contrastive structure as the explanations they serve to back.

The idea that causal claims are contrast sensitive is thus not *ad hoc* but rooted in the roles that the notion of causation plays. I should note that these role arguments might be thought to push from *Double-Contrasts* to *Total-Contrasts*, if for instance we are looking at cases involving difference-making *chains* where we need to think of the connection from c to e as mediated via d . But since I am maintaining neutrality between *Double-Contrasts* and *Total-Contrasts*, this is a matter I will leave for further discussion.

3 SEMANTICS FOR CONTRASTIVISTS?

So far I have argued that causal claims are semantically context sensitive as per *Causal Contextualism* (section 1), and that the sensitivity involved

is sensitivity to the salient contrasts to *c* and *e* as per *Double-Contrasts* or *Total-Contrasts* (section 2). This picture of semantic sensitivity to contrasts invites a natural follow-up question, as to whether and how semantic sensitivity to contrasts can be implemented within an overall plausible semantic framework. This will turn out to be something of a puzzle, with two connected aspects.

The first aspect of the puzzle concerns the existence of any semantic level parameter or parameters that tracks the kind of *bi-contrastivity* I have posited for causal claims. While there is good reason to posit a semantic level parameter (namely *the question under discussion*) which generally provides for contrasts, it is much more difficult to motivate any general provision for two separate reservoirs of contrasts (*contrasts specifically for the cause* and *contrasts specifically for the effect*). The second aspect of the puzzle concerns the connection between the semantic level parameter or parameters and the denotation of ‘causes.’ Even given a general provision for two separate reservoirs of contrasts, the clause spelling out the denotation of ‘know’ must pick up on these parameters in a precedented and plausible way.

I remain hopefully that this twofold puzzle can be solved, but cannot yet offer anything like a satisfactory solution. Accordingly I must conclude that we do not yet have a clear understanding of context sensitivity as it arises for causal claims. This is everyone’s problem. It arises in a specific form given the sort of semantic bi-contrastivity I have argued for. But the problem re-arises in different forms for different approaches. (For instance, if one thinks that there is merely pragmatic default sensitivity instead, one needs to show how this fits into an overall plausible pragmatic framework.) In this respect the context sensitivity of causal claims might be especially interesting to the student of context sensitivity, insofar as the data seems strong but the theoretical treatment difficult.

3.1 The Problem of Bi-Contrastivity

To begin with, there is good reason to posit a semantic parameter which generally provides for contrasts. This parameter is the *question under discussion*, posited as an element of the contextual scoreboard. The question under discussion (or perhaps better: a stack of questions, with the topmost element being under discussion) is widely posited to explain various phenomena such as topic choice and the licensing of ellipsis (Roberts 2004).

Questions are sets of alternatives. For instance, on the influential account of Groenendijk and Stokhof (1984), questions are *partitions* on logical space. So for instance, given that John, Billy, and Mary are the contextually salient individuals, the intensions of the question *who ate the cookies* is the set whose eight members are:

John	Billy	Mary	
Y	Y	Y	(John, Billy, and Mary all ate the cookies)
Y	Y	N	(John and Billy ate the cookies, but Mary did not)
Y	N	Y	
Y	N	N	
N	Y	Y	
N	Y	N	
N	N	Y	(Only Mary ate the cookies)
N	N	N	(No one ate the cookies)

Here is a space of contrasts, as a semantic parameter on the contextual score-card. So it seems that semantic contrastivity is linguistically plausible (in a way that attributing default and model elements is not, pending any independently attested evidence that arbitrary contexts track such information).

Moreover, the idea that the causal contrasts are coming from the question under discussion directly fits the idea that the causal inquiry controls acceptability (seen in the case where John kissed Mary, and in Mackie's account of causal selection). If the causal inquiry forms the question under discussion and thus provides contrasts, and causation is a relation involving contrasts, then there is a direct link between the causal inquiry and the contrasts involved in causation.

But there are two problems with relying on the question under discussion to furnish the contrasts. The first and most glaring problem is that—at least on the form of contrastivism I have defended—one needs more than a single source of contrasts. One needs two separate sources of contrasts, one for the contrast to the cause and another for the contrast to the effect. One needs semantic bi-contrastivity, and this goes beyond what any one question can be guaranteed to provide.

The second problem is that it is not obvious that the question under discussion will provide any contrasts *for the cause* or *for the effect*. For instance, imagine that the king has eaten soup and died, and that the question under discussion is whether there is any connection between these events. In such a context, the occurrence of both the candidate cause and effect events are presupposed. It seems as if alternatives are not being queried as to what the king ate, or what fate he suffered. The alternative being queried are of the wrong sort entirely to provide either of the contrasts needed. (This shows that the problems are not solved by trying to retreat from *Double-Contrasts* to *Cause-Contrast* or *Effect-Contrast*.)

Full disclosure: I do not know how to solve these problems. But here is one thought which may not be hopeless, which involves thinking of causal claims as obligatorily triggering a specific sort of question under discussion. There are conjunctive questions which not only provide the kind of

bi-contrastivity required, but which moreover specifically target the cause and effect slots. Most generally, such questions take the form:

26. What happened, and why?

The ‘What happened?’ aspect of 26 is understood to provide alternatives to the effect, and the ‘Why?’ aspect is understood to provide alternatives to the cause. For instance, in the case of the king just above, the true answer to “What happened?” is that the king died, and the other possible answers might include the option that the king is merely sleeping; while the true answer to “Why?” might be that the king ate the soup, and the other possible answers might be that the king poured the soup down the drain.

If it could be maintained that causal claims obligatorily trigger a question under discussion of the sort exemplified by 26, then all would be well (at least with respect to the problem of bi-contrastivity). But I should like some independent reason to maintain that causal claims obligatorily trigger such a question under discussion, beyond the fact that it would help me out. Any such triggered question ought to *show up* in topic choice and the licensing of ellipsis, in just the way that the question under discussion generally makes itself manifest in discourse (which provide the very rationale for positing the question under discussion parameter in the first place). I am not convinced that the question I am suggesting may be triggered shows up in the right ways.

3.2 The Problem of Connection

Suppose that the problem of bi-contrastivity (section 3.1) is somehow surmounted; there still remains a problem of how to connect the presence of contrasts in the context to the truth-conditions of the causal claim. After all, there are many elements of the contextual scorecard, and not every denotation is sensitive to every element. For instance, assuming that there is an element of the contextual scorecard for who is being addressed, presumably the denotation of ‘and’ remains contextually invariant, and the denotation of ‘I’ remains contextually variant but still invariant with respect to that parameter. So what is it about ‘know’ that connects it to that parameter, and thus enables it to pick up on the contextually given contrasts?

It has been argued that there are tight constraints between context sensitivity and logical form. In particular Stanley argues that all semantic context sensitivity arises from either indexicality or something like a covert variable: “Any contextual effect on truth-conditions that is not traceable to an indexical, pronoun, or demonstrative in the narrow sense must be traceable to a structural position occupied by a variable” (2000, 401). This is an attractive picture insofar as it provides principled constraints on context sensitivity, especially so given that there are

principled tests for indexicals and for covert variables. So—assuming Stanley’s constraints on semantic context sensitivity—there are three main options: either ‘cause’ is an indexical, or it projects covert contrast variables which may be evaluated by context, or it is not really semantically context sensitive after all.

I think it should be fairly clear that ‘cause’ is no indexical. Indeed it seem to fail all standard tests for indexicality.¹⁵ For instance, we automatically adjust indexicals in indirect quotation. If Ann says ‘I’m thirsty’ we report ‘Ann said that she is thirsty,’ shifting automatically from her ‘I’ to our ‘she.’ We do *not* report homophonically by ‘Ann said that I’m thirsty.’ Nothing like this seems to occur with causal claims. If the engineer concerned with the roads says:

1. The poor road conditions caused the accident

Then it seems that she may be homophonically reported in any context, even the context of the detective concerned with the drunk driver, via:

27. The engineer said that the poor road conditions caused the accident

Or at least there is nothing like the smooth and automatic adjustment of indexicals across contexts.

So can the context sensitivity of causal claims be understood in terms of a covert variable (or perhaps two covert contrast variables) instead? Perhaps so, but again I should like some independent reason to posit such variables, beyond the fact that it would help me out. Any such variables ought to *show up* in standard tests for covert variables. But—at least with respect to *the binding test* (Partee 1989; Stanley 2000)—no such variables seem to turn up:

28. *Rather than any other setting, the switch’s being set to *local* caused the passengers to arrive at the station.
29. *Rather than any other outcome, the switch’s being set to *local* caused the passengers to arrive at the station.

Perhaps there are ways to explain binding failures compatible with the presence of the covert variable. Perhaps there are other syntactic diagnostics that would render a different verdict. But *prima facie* there does not look to be a covert variable in the syntax. And so it seems that—at least if Stanley’s constraints are accepted—then there are good arguments against locating the context sensitivity of causal claims anywhere in the semantic machinery.

Since I am characterizing a parameter as semantic when it impacts truth-conditions, there remain several other options. One option would be to add contrasts parameters to the index by which propositions are

evaluated for truth, alongside the orthodox world and time parameters posited by Kaplan. But this strikes me as unpromising, since in my view it was a mistake all along to have such an index at all (Schaffer *forthcoming*). At any rate the standard reason for wanting parameters—namely the existence of a sentential operator said to work by shifting them—does not seem to apply. Another option would be to allow for *free enrichment* whereby considerations of general rationality can add constituents that lack representation in logical form (Sperber and Wilson 1986). But the worry with such an option—which is a main motivation for Stanley’s claim that “all truth-conditional effects of extra-linguistic context can be traced to logical form” (2000, 391)—is that it overgenerates, undoing needed constraints.

3.3 Concluding Puzzlement

What emerges is an inconsistent triad of seemingly plausible claims:

Pragmatic or Semantic: Causal claims are either pragmatically context sensitive or semantically context sensitive (section 1.1).

Not Pragmatic: Causal claims are not pragmatically context sensitive (section 1.3).

Not Semantic: Causal claims are not semantically context sensitive (sections 3.1–3.2).

Something must go. One either needs to reconsider all of the examples of section 1.1 so as to deny *Pragmatic or Semantic*; or one needs a better account of the pragmatic mechanisms in play, that will enable one to deny *Not Pragmatic*; or one needs a better account of how contextual elements can impact the semantics, that will enable one to deny *Not Semantic*. In other words, the context sensitivity of causal discourse seems to fit neither the Gricean view of pragmatics nor Stanley’s constraints on semantic context sensitivity.

I hold out hope that *Not Semantic* can be answered. Or at least, it seems to me that the case for *Pragmatic or Semantic* is extremely strong, turning on “data” that has been universally accepted since Mill. And it seems to me that the case for *Not Pragmatic* is fairly strong as well, at least on anything like a Gricean picture. In contrast I think that the case for *Not Semantic* is a good deal weaker, involving controversial matters concerning the question under discussion and strong views on how truth conditions are constrained by syntax. But that said, these remain serious problems that I do not know how to resolve.

So I must conclude that there is as of yet no decent account of the context sensitivity of causal claims, invariantist or contextualist. Causal context sensitivity paddles, waddles, and quacks like semantic contrast sensitivity. But where are the ducks in the semantics?¹⁶

NOTES

1. Causal selection is often assimilated to the context sensitivity of ‘the cause.’ But these phenomena should be distinguished. On the one hand, multiple factors may be selected. The engineer, for instance, might select both the presence of the potholes and the absence of a stop sign as causes of the accident. On the other hand, the context sensitivity of ‘the cause’ is at least partly a matter of the separate context sensitivity of ‘the’ and does not obviously have anything to do with ‘cause’ (any more than the context sensitivity of ‘the dog’ automatically establishes any context sensitivity for ‘dog’).
2. Anscombe (1969) provides a similar example. She notes that one might accept that de Gaulle’s making a speech caused an international crisis, but deny that the man with the biggest nose in France’s making a speech caused an international crisis (without denying the facts). One wants to say: the size of the nose was not relevant. But unless one has an implausibly fine conception of events, there was only one speech.
3. While it is theoretically possible to reject both *Causal Contextualism* and *Causal Invariantism* (for instance by rejecting the “data” of the previous section), I am not aware of any theorists who have taken this approach. One theoretically alternative that does come up is to treat ‘cause’ as semantically *ambiguous*, as per Davidson’s suggestion that ‘cause’ is ambiguous between the relation of causation and the sentential connective of causal explanation (1980, 162). But—though I lack the space for a proper discussion of Davidson’s suggestion—I do not think it withstands much scrutiny. To my knowledge no serious linguistic evidence for any such ambiguity has been mooted, nor is there any reason to think that the sensitivities illustrated are due to “disambiguation.” Indeed the ambiguity claim should entail that all of the causal claims considered have true readings, which should thereby be favored by charity. But the data is rather that seemingly “true” causal claims—such as the claim in 3 that the presence of oxygen caused the forest fire—still count as *unacceptable* in certain contexts. So unless some interpretive pressures are revealed which might overturn charity, ambiguity claims just get the data wrong. (In general, ambiguity claims multiply opportunities to find acceptable interpretations, and so they are good for explaining acceptabilities but not so good for explaining unacceptabilities.)
4. Lewis does tolerate some semantic context sensitivity in causal discourse, both with respect to the vagueness of counterfactuals (“The vagueness of similarity does infect causation, and no correct analysis can deny it” (1986, 163)) and—in his later influence account (2000)—with respect to the degree of influence sufficient for counting as a cause. But these look to be independent from the context sensitivities in section 1.1. So Lewis is perhaps best classified as a friend of *Causal Contextualism*, but one who sides with *Causal Invariantism* with respect to the issues under consideration in the main text.
5. I work with the Gricean view of conversational pragmatics simply because it is the most orthodox and developed approach. The invariantist who prefers a different view of pragmatics should take the discussion in the main text as a challenge to do better.
6. For discussions connecting the Gricean maxim of Relevance to the question under discussion (“speak to the question!”) see Roberts 2004.
7. Here I am generalizing an argument due to McGrath, who considers pragmatic explanations for why we deny that certain omissions are causes (for instance, why we deny that the queen of England’s failing to water my

flowers caused them to wilt), and points out: “it isn’t just that we refuse to utter [omission sentences] that are, on the view, true; we also utter their negations” (2005, 128–129). Similar points arise in the literature on epistemic contextualism, as brought out by DeRose (1999).

8. Other advocates of the view that causal reasoning involves the notion of deviation from a default include Maudlin (2004) and McGrath (2005). But for Maudlin the notion of a default is encoded in the laws of nature (to the extent it is recoverable at all), and for McGrath the notion of a default comes from our notion of what is normal. As far as I can see, neither explicitly allows for context sensitivity, although both certainly could. Indeed McGrath’s notion of what is normal strikes me as most naturally understood as a context-sensitive notion. What counts as “normal” for the forest rangers may be quite different from what counts as “normal” for the visiting Venusians.
9. This is perhaps clearest in Halpern’s (2000) formalism for causal models, in which one begins from a *signature* $\langle U, V, R \rangle$, where U is a set of *exogenous variables* (“initial conditions of the system”), V is a set of *endogenous variables* (“dependent conditions”), and R is a *range function* associated each variable $X \in U \cup V$ with a range of at least two allotted values. R encodes contrasts for the totality of events represented in the model. See also Eagle 2007 and Schaffer 2010.
10. Hall (2007) uses defaults to distinguish two sorts of causal structures which standard Pearl models conflate. To my mind this is the most promising case to be made for thinking that the notion of default is also essential to characterizing causal notions.
11. Arguably analogues of all of these sensitivities are to be found in knowledge ascriptions (Schaffer 2005b), with the seeming exception of selection effects. There does not seem to be anything on the epistemic side corresponding to selection (this would be a contextually variable tendency to promote certain elements from a subject’s body of knowledge to the level of knowledge, and to demote the remaining elements to background information). This seeming disanalogy should be a mystery for everyone. For those who go in for parallel treatments of the context sensitivity of causal claims and knowledge ascriptions (such as myself), the mystery is why there is a minor *disanalogy*. For those who do not go in for parallel treatments, the mystery is why there is a *minor* disanalogy. Since my treatments of the context sensitivity of causal claims and knowledge ascriptions involve some minor differences (with knowledge claims there is only one contrast, with causal claims there are two), I should like to appeal to these minor differences to explain the minor disanalogy, but I know not how.
12. More precisely, Rooth adds a semantic focus marker whose value is a contextually determined set of options, and posits a dual interpretation of phrases with this marker. To illustrate, ‘Socrates’s DRINKING HEMLOCK at dusk’ is semantically interpreted as $[\dots][\text{Socrates’s}[\text{drinking hemlock}]_F \text{ at dusk}]\dots$, where $[\text{drinking hemlock}]_F$ induces a dual interpretation: there is the “ordinary semantic value” of drinking hemlock, and the “focus semantic value” which is the set of contextually salient options for what Socrates might have done at dusk (including drinking hemlock, but including other options as well). Semantic sensitivity to focus is then understood in terms of operators sensitive to these focus semantic values. Given that causal claims exhibit semantic sensitivity to focus, and given Rooth’s alternative semantics for focus, it falls out that ‘cause’ is contrast sensitive.
13. *Lacuna*: if 3 does not receive any natural interpretation than its denial should not either, which does not quite fit that data in 13. So it would be smoother for me to say that 3 does receive some interpretation as a contrastive falsehood in

the context of the forest rangers. But I do not currently have any contrastive falsehood to suggest for the role.

14. Selection is the one aspect of context sensitivity that seems not to apply equally to both the cause and effect side, operating primarily on the cause side. There may also be something like selection on the effect side in our intuitive distinction between causes and byproducts, but this matter needs further exploration. (Selection seems to have a variety of special features, and may ultimately need separate treatment).
15. Cappelen and Lepore (2005) provide a useful battery of tests for indexicality (I think they mistake these for tests for context sensitivity generally, and don't properly consider the prospect that context sensitivity might come in multiple forms.)
16. Thanks to Mark Heller, Chris Hitchcock, Cei Maslen, and Peter Menzies.

REFERENCES

- Achinstein, P. 1975. Causation, Transparency, and Emphasis. *Canadian Journal of Philosophy* 5: 1–23.
- Anscombe, G.E.M. 1969. Causality and Extensionality. *Journal of Philosophy* 66: 152–159.
- Bennett, J. 1995. *The Act Itself*. Oxford: Oxford University Press.
- Cappelen, H., and E. Lepore. 2005. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Basil Blackwell.
- Davidson, D. 1980. Causal Relations. In *Essays on Actions and Events*, 149–162. Oxford: Oxford University Press.
- DeRose, K. 1999. Contextualism: An Explanation and Defense. In *The Blackwell Guide to Epistemology*, ed. J. Greco and E. Sosa, 187–205. Oxford: Basil Blackwell.
- Eagle, A. 2007. Pragmatic Causation. In *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*, ed. H. Price and R. Corry, 156–90. Oxford: Oxford University Press.
- Garfinkel, A. 1981. *Forms of Explanation: Rethinking the Questions in Social Theory*. New Haven: Yale University Press.
- Groenendijk, J., and M. Stokhof. 1984. Questions. In *Handbook of Logic and Language*, ed. J. van Benthem and A. ter Meulen, 1055–1124. Amsterdam: Elsevier.
- Hall, N. 2007. Structural Equations and Causation. *Philosophical Studies* 132: 109–136.
- Halpern, J. 2000. Axiomatizing Causal Reasoning. *Journal of Artificial Intelligence Research* 12: 317–337.
- Halpern, J.Y., and J. Pearl. 2005. Causes and Explanations: A Structural-Model Approach. Part 1: Causes. *British Journal for the Philosophy of Science* 56: 843–887.
- Hart, H.L.A., and A.M. Honoré. 1985. *Causation in the Law*. Oxford: Oxford University Press.
- Hitchcock, C. 1996. The Role of Contrast in Causal and Explanatory Claims. *Synthese* 107: 95–419.
- Hitchcock, C. 2003. Of Humean Bondage. *British Journal for the Philosophy of Science* 54: 1–25.
- Hitchcock, C. 2007. Prevention, Preemption, and the Principle of Sufficient Reason. *Philosophical Review* 116: 495–532.
- Kierkegaard, S. 1978. *Parables of Kierkegaard*, ed. T.C. Oden. Princeton: Princeton University Press.

- Lewis, D. 1986. *Philosophical Papers*. Vol. 2. Oxford: Oxford University Press.
- Lewis, D. 2000. Causation as Influence. *Journal of Philosophy* 97: 182–197.
- Mackie, J.L. 1974. *The Cement of the Universe*. Oxford: Oxford University Press.
- Maslen, C. 2004. Causes, Contrasts, and the Nontransitivity of Causation. In *Causation and Counterfactuals*, ed. J. Collins, N. Hall, and L.A. Paul, 341–357. Cambridge, MA: MIT Press.
- Maudlin, T. 2004. Causation, Counterfactuals, and the Third Factor. In *Causation and Counterfactuals*, ed. J. Collins, N. Hall, and L.A. Paul, 419–443. Cambridge, MA: MIT Press.
- McDermott, M. 1995. Redundant Causation. *British Journal for the Philosophy of Science* 40: 523–544.
- McGrath, S. 2005. Causation by Omission: A Dilemma. *Philosophical Studies* 123: 125–148.
- Menzies, P. 2004. Difference-Making in Context. In *Causation and Counterfactuals*, ed. J. Collins, N. Hall, and L.A. Paul, 139–180. Cambridge, MA: MIT Press.
- Menzies, P. 2007. Causation in Context. In *Causation, Physics, and the Constitution of Reality*, ed. H. Price and R. Corry, 191–223. Oxford: Oxford University Press.
- Mill, J.S. 1950. *A System of Logic*. New York: Macmillan.
- Partee, B. 1989. Binding Implicit Variables in Quantified Contexts. *Proceedings of the Chicago Linguistics Society* 25: 342–365.
- Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.
- Putnam, H. 1982. Why There Isn't a Ready-Made World. *Synthese* 51: 141–167.
- Roberts, C. 2004. Context in Dynamic Interpretation. In *The Handbook of Pragmatics*, ed. L. Horn and G. Ward, 197–220. Oxford: Basil Blackwell.
- Rooth, M. 1992. A Theory of Focus Interpretation. *Natural Language Semantics* 1: 75–116.
- Schaffer, J. 2005a. Contrastive Causation. *Philosophical Review* 114: 327–358.
- Schaffer, J. 2005b. Contrastive Knowledge. *Oxford Studies in Epistemology* 1: 235–271.
- Schaffer, J. 2010. Contrastive Causation in the Law. *Legal Theory* 16: 259–297.
- Schaffer, J. *forthcoming*. Necessitarian Propositions. *Synthese*.
- Sperber, D. and D. Wilson. 1986. *Relevance: Communication and Cognition*. Oxford: Basil Blackwell.
- Stanley, J. 2000. Context and Logical Form. *Linguistics and Philosophy* 23: 391–434.
- Van Fraassen, B. 1980. *The Scientific Image*. Oxford: Oxford University Press.
- Woodward, J. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.