

JONATHAN SCHAFFER

## THE IRRELEVANCE OF THE SUBJECT: AGAINST SUBJECT-SENSITIVE INVARIANTISM

“The death of man is nothing to get particularly excited about. It’s one of the visible forms of a much more general decease, if you like. I don’t mean by it the death of god but the death of the subject, of the Subject in capital letters, of the subject as origin and foundation of Knowledge, of Liberty, of Language and History.” – Michel Foucault

Does what you know depend on what is at stake for you? That is, is the knowledge relation sensitive to the subject’s practical interests? Subject sensitive invariantists (Fantl and McGrath, 2002; Hawthorne, 2004, ch. 4; Stanley, forthcoming) say that the answer is *yes*. They claim to capture the contextualist data without the shifty semantics. I will argue that the answer is *no*. The knowledge relation is sensitive to what is in question for the attributor, rather than what is at stake for the subject. There is no substitute for the contextualist semantics.

### 1. STAKES SHIFTING INTUITIONS

According to Subject Sensitive Invariantism (henceforth: SSI), what you know depends on what is at stake for you. SSI is advertised as *the best explanation for stakes shifting intuitions*. So I will begin by examining these intuitions.

The intuitions in question arise in pairs of cases such as the following:

*Low*: Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks, though it is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the

lines inside are very long, as they often are on Friday afternoons. Realizing that it isn't very important that their paychecks are deposited right away, Hannah says "I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning."

*High*: Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their checking account, it is very important that they deposit their paychecks by Saturday. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, "I guess you're right. I don't know that the bank will be open tomorrow." (Stanley, forthcoming, p. 3)

Most people intuit that Hannah knows in *Low* but does not know in *High*. Cohen (1988, 1999) and DeRose (1992, 1995) use this sort of intuitive difference to argue for contextualism. Fantl and McGrath (2002) and Stanley (forthcoming) offer SSI as an invariantist proposal that fits the contextualist data, since what is at stake for the subject differs in *Low* and *High*.<sup>1</sup>

Since cases such as *Low* and *High* provide the main motivation for SSI, it will prove important to attend to the details. Notice first that *Low* and *High* differ in many ways. They differ not only over what is at stake for the subject, but also over whether the possibility of error is explicitly mentioned (which only happens in *High*).<sup>2</sup> As such *Low* and *High* cannot show *which difference* is driving our intuitions. Notice second that *Low* and *High* explicitly include the SSI-ers preferred verdicts. Hannah asserts what the SSI-er hopes the reader will intuit. As such, *Low* and *High* are presented in a biasing manner.

To test whether our intuitions are sensitive to what is at stake for the subject, one needs *unbiased minimal pairs* – cases differing only over what is at stake for the subject, without telling the reader what she is supposed to intuit. So here are better test cases:

*Low-stakes*: On Friday afternoon, Sam is driving past the bank with his paycheck in his pocket. The lines are long. Sam would prefer to deposit

his check before Monday, but he has no pressing need to deposit the check. He has little at stake. Sam remembers that the bank was open last Saturday, so he figures that the bank will be open this Saturday. He is right – the bank will be open. So, does Sam know that the bank will be open this Saturday?

*High-stakes:* On Friday afternoon, Sam is driving past the bank with his paycheck in his pocket. The lines are long. Sam would prefer to deposit his check before Monday, and indeed he has pressing financial obligations that require a deposit before Monday. His entire financial future is at stake. Sam remembers that the bank was open last Saturday, so he figures that the bank will be open this Saturday. He is right – the bank will be open. So, does Sam know that the bank will be open this Saturday?

Here I tend to intuit that Sam knows, in both cases. Or at least, the intuitions are nowhere near as clear. If that is right, then something else – not the stakes – must be driving our intuitions in *Low* and *High*. Perhaps that something else is the attributor's consideration of error.

To test whether attributor considerations of error are what really drive our intuitions in *Low* and *High*, one should consider unbiased minimal pairs that differ only over the salience of error. Thus one should compare *Low-stakes* to:

*Low-and-Salient:* On Friday afternoon, Sam is driving past the bank with his paycheck in his pocket. The lines are long. Sam would prefer to deposit his check before Monday, but he has no pressing need to deposit the check. He has little at stake. Sam remembers that the bank was open last Saturday, so he figures that the bank will be open this Saturday. He is right – the bank will be open. But banks do change their hours, and Sam has not looked into this. So, does Sam know that the bank will be open this Saturday?

And one should compare *High-stakes* to:

*High-and-Salient:* On Friday afternoon, Sam is driving past the bank with his paycheck in his pocket. The lines are long. Sam would prefer to deposit his check before Monday, and indeed he has pressing financial obligations that require a deposit before Monday. His entire financial future is at stake. Sam remembers that the bank was open last Saturday, so he figures that the bank will be open this Saturday. He is right – the bank will be open. But banks do change their hours, and Sam has not looked into this. So, does Sam know that the bank will be open this Saturday?

For both *Low-and-Salient* and *High-and-Salient*, I tend to intuit that Sam does not know. So overall, when the possibility of error is not made salient (*Low-stakes* and *High-stakes*), I tend to intuit that the subject knows; but when the possibility of error is made salient (*Low-and-Salient* and *High-and-Salient*), I tend to intuit that the subject does not know. The stakes play no role. Perhaps my intuitions are unusual, and no doubt they are theoretically biased. But it seems to me at least, that when one presents the cases carefully, in ways that actually *discriminate* between the SSI explanation and the contextualist explanation, only the contextualist explanation fits the data.

In any case, even if one still feels the pull of the stakes when the cases are presented carefully, I submit that any such intuitions are easily reversed. So consider:

*Low-and-Slow*: On Friday afternoon, Sam is driving past the bank with his paycheck in his pocket. The lines are long. Sam would prefer to deposit his check before Monday, but he has no pressing need to deposit the check. He has little at stake. Sam remembers that the bank was open last Saturday, so he figures that the bank will be open this Saturday. He is right – the bank will be open.

As Sam is about to drive on, his car dies, right beside the bank. Now he has an hour to kill before the tow truck comes. He could easily deposit his check, or at least look at the hours posted on the door to confirm that the bank will be open this Saturday. But instead Sam just dozes in the backseat. So, does Sam know that the bank will be open this Saturday?

*High-and-Fast*: On Friday afternoon, Sam is driving past the bank with his paycheck in his pocket. The lines are long. Sam would prefer to deposit his check before Monday, and indeed he has pressing financial obligations that require a deposit before Monday. His entire financial future is at stake. Sam remembers that the bank was open last Saturday, so he figures that the bank will be open this Saturday. He is right – the bank will be open.

As Sam is about to stop to double-check the bank hours, he remembers that he promised to buy a present for his wife. She will be furious if he forgets – his whole relationship is at stake. The stores are about to close. Sam must choose. So Sam makes a split-second decision to drive past the bank and pick up a present for his wife instead, thinking that after all, the bank will be open this Saturday. So, does Sam know that the bank will be open this Saturday?

I tend to intuit that Sam does not know in *Low-and-Slow*, but does know in *High-and-Fast*. In *Low-and-Slow*, it seems to me that Sam does not know since he should have double-checked. He had all the time in the world. He was epistemically negligent. While in *High-and-Fast*, it seems to me that Sam does know since he made the informed choice of getting the present for his wife. He knew enough to choose wisely. At any rate, I doubt anyone will intuit that the subject knows in *Low-and-Slow* but does not know in *High-and-Fast* – which is what SSI predicts.<sup>3</sup> This is further evidence that it is not the stakes that are driving our intuitions. Perhaps what is driving our intuitions in *Low-and-Slow* is the thought that Sam should have double-checked, which is suggestive to us of the possibility that Sam might be in error.

So I conclude that the core motivation for SSI does not stand scrutiny. Set up the cases carefully – as unbiased minimal pairs – and you’ll see that what drives our intuitions is not what is at stake for the subject, but rather what is salient for the attributor. Think what you will about the contextualist semantics. But only contextualism is capturing the data.<sup>4</sup>

## 2. ANXIOUS ATTRIBUTORS

Whatever the fate of the stakes shifting intuitions, the SSI semantics still does not account for other intuitions in the neighborhood. In particular, as Hawthorne and Stanley both explicitly acknowledge, the SSI semantics goes wrong in the following sort of case:

*High-on-Low*: On Friday afternoon, Sam is driving past the bank with his paycheck in his pocket. The lines are long. Sam would prefer to deposit his check before Monday, but he has no pressing need to deposit the check. He has little at stake. Sam remembers that the bank was open last Saturday, so he figures that the bank will be open this Saturday. He is right – the bank will be open.

You, by the way, have your entire financial future at stake here. If Sam doesn’t deposit his check before Monday, Sam’s check to you will not clear in time to save you from impending bankruptcy. Sam has not

bothered to look into whether the bank might have changed its hours. So, does Sam know that the bank will be open this Saturday?

Here the SSI semantics has it that Sam knows (since little is at stake *for him*), but intuitively Sam does not know. Never mind whether what drives us to impute ignorance to Sam is our high stakes, or our anxiety-driven consideration of possibilities of error. Both are present in *High-on-Low*. Either way, the SSI semantics gets the case wrong.

Hawthorne and Stanley thus propose several pragmatic and error-theoretic explanations, which I will now consider. Hawthorne's first proposal (2004, p. 160) is that the high-stakes attributor is *pragmatically unable to assert that Sam knows*. Given that one must know that *p* to assert that *p* (Williamson, 2000, ch. 11), the attributor (you) must know that Sam knows, to assert that Sam knows. Given SSI, and your high stakes situation in *High-on-Low*, you will presumably fail to know this. Thus Hawthorne concludes that SSI plus the knowledge account of assertion explains the data here. But first, this can only explain why the high-stakes attributor will not assert that Sam knows – it will not explain why the high-stakes attributor will go so far as to assert that Sam does not know (Stanley, forthcoming, pp. 75–76). And second, the explanation only works if the high-stakes attributor lacks the evidence needed to overcome her high stakes (DeRose, forthcoming). In *High-on-Low*, I have stipulated that the bank indeed will be open, which should assuage any stakes-related anxieties for you. You have all the evidence you could want. You know the bank will be open. You are just being critical of Sam.<sup>5</sup>

Hawthorne's second proposal (2004, p. 164) is that the high-stakes attributor makes a *projection error*. We are generally prone to project our anxieties, so perhaps the attributor is projecting her anxieties onto the low-stakes subject and so treating her as if she faced high stakes. Given SSI, that would explain why we speak of the low-stakes subject as ignorant. But first, this proposal overgenerates: the sort of projection error Hawthorne has in mind would equally explain away

any residual intuition we have to deny knowledge to the subject in the original *High*. Perhaps the high-stakes subject does know, but we intuit otherwise because we overestimate the possibility that he is at risk. This would thus undermine the original case for SSI (Williamson, forthcoming). Second, the projection error line is inapplicable to the case as described. In the case as described it is quite clear that only you are at risk of bankruptcy. Sam is safe. We can test for this by asking *both* whether Sam knows *and* whether Sam is at risk. If you answer that Sam does not know *and* is not at risk, then your basis for denying him knowledge cannot be the sort of projection error Hawthorne posits.

Stanley's proposal (forthcoming, pp. 77–79) is that the high stakes attributor makes a *semantic error*. Perhaps we fixate on our own situation, and so tend to confuse the proposition that the subject knows, with the counterfactual proposition that if the subject were in our own situation, then she would know. If so, given SSI, we would confuse the true proposition that the low-stakes subject knows, with the false counterfactual proposition that if the subject were in my high-stakes situation, then she would not know. But first, the explanation wrongly predicts that we would confusedly attribute knowledge in *High*, by fixating on our own low-stakes situation there. The sort of semantic confusion Stanley has in mind would equally predict that the low-stakes attributor should ascribe knowledge to the high-stakes subject. If we are prone to such a confusion, we should be equally prone to confuse the (false, given SSI) proposition that the high-stakes subject knows, with the (true, given SSI) proposition that if the high-stakes subject were in the attributor's low-stakes situation, then the subject would know. This would thus undermine the original case for SSI. And second, I see no precedent for imputing such an error. I am not at all inclined, for instance, to think that Sam must be anxious about his financial future, because I confuse that (false) proposition with the (true) counterfactual proposition that if Sam were in my situation, then he would be anxious about his financial future. At least,

some independent evidence is needed that we actually are prone to such confusions.

Perhaps the SSI-er can offer some further proposals. Perhaps other theories require equally implausible pragmatic and/or error-theoretic lines. In any case, it is worth noting that the SSI-ers are not avoiding pragmatic and error-theoretic explanations. There is no purely semantic, invariantist explanation to be had here. Indeed, the SSI-ers are not even garnering improved pragmatic and/or error-theoretic explanations at this point. *High-on-Low* thus poses a particularly bad problem for SSI – it shows that SSI is not even outpacing the competition in the very cases it was designed for.

It is worth pausing here to compare how the main competitors are faring. I will focus on dogmatism, skepticism, contextualism (of a certain sort: see below), and SSI. I will give two tables of results, the first for the intuitions as the SSI-ers report them, and the second for the intuitions as I report them above. Table 1 (below) shows the results for the intuitions as the SSI-ers report them. So by the SSI-ers own lights, the SSI semantics is not outpacing either the skeptical or the contextualist competition. And Table 2 (below) shows the results for the intuitions as I report them. So by my lights, the SSI semantics comes in *last place*, in the very cases it was designed for. Of course, even so unsuccessful a semantics should not yet be eliminated. It is *tolerable* if a semantic theory gets some black marks, even in the very cases it was designed for, *if* the theory can be supplemented with a decent pragmatic or error-theoretic explanation. But so far the prospects seem bleak.<sup>6</sup> Especially when the theory must compete with a runaway winner: contextualism.

Table 1. Intuitions as SSI-ers report them

	Dogmatism	Skepticism	Contextualism	SSI
<i>Low-stakes</i> (knowledge)	√	X	√	√
<i>High-stakes</i> (ignorance)	X	√	X	√
<i>High-on-Low</i> (ignorance)	X	√	√	X
<i>High-on-High</i> (ignorance)	X	√	√	√



Table 2. Intuitions as I would report them

	Dogmatism	Skepticism	Contextualism	SSI
<i>Low-stakes</i> (knowledge)	√	X	√	√
<i>High-stakes</i> (knowledge)	√	X	√	X
<i>Low-and-Salient</i> (ignorance)	X	√	√	X
<i>High-and-Salient</i> (ignorance)	X	√	√	√
<i>Low-and-Slow</i> (ignorance)	X	√	√	X
<i>High-and-Fast</i> (knowledge)	√	X	√	X
<i>High-on-Low</i> (ignorance)	X	√	√	X

### 3. THE SOCIAL ROLE OF KNOWLEDGE

Things get worse for SSI. I will now argue that, single case intuitions aside, there are general theoretical reasons for not rendering knowledge sensitive to what is at stake for the subject, stemming from the social role of knowledge.

To begin with, knowledge has a social role, linked to inquiry, expertise, testimony, norms of conversation, and what we value, *inter alia*. Here are some very rough approximations of plausible principles connecting these notions:

*Knowledge closes inquiry*: if there is an inquiry as to whether  $p$ , and  $S$  comes to know that  $p$ , then  $S$  is in position to close the inquiry ( $S$  knows the answer).

*Knowledge grounds expertise*: if, for all propositions  $p$  on topic  $t$ ,  $S$  knows that  $p$ , then  $S$  is an expert on  $t$  ( $S$  knows everything about  $t$ ).

*Testimony transmits knowledge*: if  $S$  knows that  $p$ , and testifies as much, then  $S$ 's hearers are in position to know that  $p$  on the basis of  $S$ 's testimony ( $S$  has informed the audience).

*Questioning and answering require knowledge*: one should ask someone a question only if she knows the answer (only ask those who know); one should answer a question only if one knows the answer (only answer if you know).

*Knowledge has value*: if ( $p$  is true,  $S$  does not know that  $p$ , if  $C$  were to obtain then  $S$  would come to know that  $p$ , and  $S$  can make  $C$  obtain), then  $S$  ought to make  $C$  obtain *ceteris paribus* ( $S$  ought to strive for knowledge).

Again, these are very rough approximations – in some places I'm not sure how to be more precise, in other places I just

want to avoid distracting details – but I hope these are at least *prima facie* plausible. For I will now argue that in order to play *anything like these roles*, the knowledge relation must be *insensitive* to what is at stake for the subject.

Starting with the connection to inquiry, imagine that Holmes and Watson are jointly inquiring into who murdered Smith. Holmes is next on the murderer's list and Watson is not, so there is a lot more at stake for Holmes. The duo finds circumstantial evidence that Black did the deed – or at any rate, the duo finds what the SSI-er will count as sufficient evidence for Watson given his low stakes, but insufficient evidence for Holmes given his high stakes. So by SSI, only Watson knows that Black murdered Smith. By the principle that knowledge closes inquiry, Watson is in a position to close the inquiry. His low stakes have given him a competitive edge. But surely it would be wrong here to say that only Watson is able to solve the mystery.

Or imagine that Beaker and Dr. Bunsen Honeydew are chemists observing the same experiment (putting the same question to nature, as it were). Here Beaker is not so interested in the outcome – little is at stake for him. But the good doctor has staked his entire professional reputation, such as it is, on the outcome – much is at stake for him. Surely Beaker and the doctor have equal opportunity to know how the experiment comes out. Just because the doctor has more at stake than Beaker does, does not mean that it is harder for the doctor to see what results. One cannot gain a competitive advantage in scientific inquiry, for instance, by not caring about the result. So I conclude that progress in inquiry is not sensitive to what is at stake for the subject.

Turning to expertise, imagine that Thumb could not care less about birds. But Thumb is really bored, so he winds up thumbing through a few field guides. He picks up just enough information, and has just low enough stakes, so that the SSI-er will count Thumb as knowing all sorts of propositions about sparrows. So by the principle that knowledge grounds expertise, Thumb will now count as an expert on sparrows.

But of course he isn't. To see this, imagine that Wing is a trained ornithologist, who finds herself in such a high stakes situation that the SSI-er will not count Wing as knowing many propositions about sparrows. Should Thumb defer to Wing, or Wing to Thumb? By the lights of SSI, Thumb is the expert so Wing should defer. But clearly it is Thumb who should defer to Wing on the subject of sparrows.

In general, the social role of the expert is to serve as a *reservoir* of knowledge.<sup>7</sup> This requires a stability in one's pool of knowledge that is not compatible with SSI. The social status of expertise cannot fluctuate as the stakes rise and fall. For instance, one cannot gain in expertise by suddenly not caring about the topic. So I conclude that the status of expertise is not sensitive to what is at stake for the subject.

Moving to testimony, imagine that it is crucial to you that I be in my office tomorrow (perhaps you need me there to hatch your nefarious plans), but not crucial to me. I tell you that I will be in my office tomorrow. Suppose that my evidence for saying that I will be in my office tomorrow is that I've made some tentative plans to meet with a student – or at any rate, suppose that my evidence is such that the SSI-er will consider it sufficient only for low-stakes knowledge. Now by SSI, since the stakes for me are low, I know that I will be in my office tomorrow. By the principle that testimony transmits knowledge, you are in position to know that I will be in my office tomorrow on the basis of my testimony. But that means that you are in a position to gain high-stakes knowledge off my low-stakes evidence. The SSI-er should regard that as epistemic cheating (what MacFarlane (2005) dubs 'knowledge laundering'). So the SSI-er must deny that my knowledgeable testimony has the power to inform you here.

In general, the social role of testimony is to allow for the spread of knowledge. This practice proceeds without needing to know what may be at stake for the testifier. If testimony could only transmit knowledge to subjects with comparable stakes (or at least to hearers with no more at stake than testifiers), then our practice of relying on testimony would be troubled. Whether we could learn from our teachers would

depend on the details of their personal lives. So I conclude that epistemic reliance on testimony proceeds without consideration of what is at stake for the testifier.

Switching over to questioning and answering, a plausible norm for questioning is that one should ask someone a question only if she knows the answer. For instance, if someone were to ask you what I like to eat for breakfast, you would probably reply with a slightly outraged “*I don’t know – why are you asking me?*” Your sense of outrage is due to your sense of being an improper target for that question.<sup>8</sup> But such a norm for questioning requires that knowledge not be sensitive to what is at stake for the subject. I might, for instance, email several of my colleagues a question about Leibniz, without considering what might be at stake for them, and without knowledge of what is going on in their personal lives that day. I ask simply because I think they might *have the information*. As Cohen (personal communication) and DeRose (2004) have pointed out, we recruit informants because we think they might have the information we need – we don’t need to know what is at stake for them. So I conclude that the norm of questioning is not sensitive to what is at stake for the subject.

A plausible norm for answering, which parallels the norm for questioning, is that one should answer a question only if one knows the answer.<sup>9</sup> For instance, if you were to answer that I like to eat bananas for breakfast, you would not know what you were talking about (even though you would be right). To answer a question is to select the true answer from a range of alternatives. All well-formed questions are multiple-choice questions.<sup>10</sup> If you can eliminate all but one answer on a multiple-choice exam question, then you are in an ideal position to answer the question. This is as true on a practice exam where nothing is at stake, as it is on a real exam where one’s academic future may be at stake. So I conclude that the norm of answering is not sensitive to what is at stake for the subject.

The norm of answering, by the way, just is the norm for assertion. All assertions are answers. As Stalnaker (1999a, b) suggests, a context can be modeled by a set of alternatives – those options that are the ‘live options’ for the conversants.

A context so modeled just is a multiple-choice slate. Now the role of an assertion, on Stalnaker's treatment, is to reduce the context set down to the content of what one asserts. That's what an assertion does. An assertion so understood is an (at least partial) answer. It functions as an answer to the question of which 'live option' in the context set is true. This shows that SSI does not fit the knowledge account of assertion. One is in an epistemic position to assert when one is in an epistemic position to reduce the context set down to the content of what one asserts. One gets in such a position by being able to eliminate the other salient alternatives, regardless of what one has at stake. If one is a contestant on a game show like *Who Wants to be a Millionaire?* and one can eliminate all but one answer to a question, one is in position to assert the answer. This is as true when a hundred dollars are at stake, as it is when a million dollars are at stake.<sup>11</sup>

Finishing with the idea the knowledge is valuable, imagine that Floyd is in a high-stakes situation with respect to *p*, but has only low-stakes evidence for *p*. Suppose also that *p* is true, that by SSI Floyd would know that *p* if the stakes were lower, and that Floyd can make the stakes go lower for him by smoking a pound of pot. Then given the value of knowledge, Floyd ought to smoke a pound of pot (*ceteris paribus*). It becomes hard to see how knowledge as SSI renders it could be valuable. The SSI-ers conception of knowledge seems to punish the passionate inquirer, and reward the comfortably numb.

What is valuable about knowledge seems connected to the power that knowledge provides for answering our questions (Castañeda, 1980; Hookway, 1996). That's what knowledge buys. What we should strive for epistemically is to be better at inquiry, rather than numb to the results. So I conclude that SSI does not fit the value of knowledge.<sup>12</sup>

There is a moral emerging. At the core of the social role of knowledge is the idea of answering questions. Progress in inquiry is answering the question. An expert is a person with the answers. Testimony is the transmission of

answers. An assertion is an answer. The value of knowledge is the power to answer. What emerges is that the social role of knowledge ascriptions is to identify people who can help us answer our questions. In short, knowledge ascriptions finger answerers (Hookway, 1996; Schaffer, 2005a). This is the deep reason why knowledge must not be sensitive to what is at stake for the subject, but must rather be sensitive to what is in question for the attributor. Only contextualism is able to provide a socially fitting conception of knowledge.

#### 4. QUESTION SENSITIVITY

SSI was advertised as capturing the contextualist data without the shifty semantics. I will conclude by arguing that the contextualists so far have actually understated their case – the best contextualist data has nothing to do with what is at stake for the subject, but has to do with what is in question for the attributor.

To begin with, it is strange that the stakes-shifting cases ever became the flagship cases for contextualism. In no contextualist theory to date do the stakes actually play any theoretical role! Rather, what shifts, at least according to the sort of theory that I like, is a set of relevant alternatives: *what is in question*.<sup>13</sup> In any case, the flagship cases for contextualism should involve shifts in what the theory says shifts.

So consider the following cases, set up as unbiased minimal pairs differing only over the contrasts salient to the attributor:

*Who/what:*

- (a) Mary has stolen the bicycle from the toy store. The detective finds Mary's fingerprints at the scene. Does the detective know who stole the bicycle?
- (b) Mary has stolen the bicycle from the toy store. The detective finds Mary's fingerprints at the scene. Does the detective know what Mary stole?

*Whether:*

- (a) Mary has stolen the bicycle from the toy store. The detective finds Mary's fingerprints at the scene. Does the detective know whether Mary or Peter stole the bicycle?
- (b) Mary has stolen the bicycle from the toy store. The detective finds Mary's fingerprints at the scene. Does the detective know whether Mary stole the bicycle or the wagon?

*Rather:*

- (a) Mary has stolen the bicycle from the toy store. The detective finds Mary's fingerprints at the scene. Does the detective know that Mary rather than Peter stole the bicycle?
- (b) Mary has stolen the bicycle from the toy store. The detective finds Mary's fingerprints at the scene. Does the detective know that Mary stole the bicycle rather than the wagon?

*Cleft:*

- (a) Mary has stolen the bicycle from the toy store. The detective finds Mary's fingerprints at the scene. Does the detective know that it was Mary who stole the bicycle?
- (b) Mary has stolen the bicycle from the toy store. The detective finds Mary's fingerprints at the scene. Does the detective know that it was a bicycle that Mary stole?

*Focus:*

- (a) Mary has stolen the bicycle from the toy store. The detective finds Mary's fingerprints at the scene. Does the detective know that *Mary* stole the bicycle?
- (b) Mary has stolen the bicycle from the toy store. The detective finds Mary's fingerprints at the scene. Does the detective know that Mary stole *the bicycle*?

*Presupposition:*

- (a) *Someone* has stolen the bicycle from the toy store. The detective finds Mary's fingerprints at the scene. Does the detective know that Mary stole the bicycle?
- (b) Mary has stolen *something* from the toy store. The detective finds Mary's fingerprints at the scene. Does the detective know that Mary stole the bicycle?

If your intuitions are like mine, you will say that the detective knows in all the (a) cases, but does not know in any of the (b) cases. There is a uniform explanation for these intuitions. Questions, rather than-clauses, clefts, focus, and pre-

suppositions are all linguistic mechanisms for encoding contrasts.<sup>14</sup> The strongest contextualist intuitions, then, are intuitions about the *question sensitivity* (or contrast sensitivity) of knowledge.<sup>15</sup>

In this respect the contextualist ought to complain that the traditional epistemologist's fetish for knowledge-*that* constructions has biased the debate. Knowledge-*wh* constructions such as I know what time it is, I know who is coming to dinner, and I know why the caged bird sings, are actually more natural in ordinary language, and more explicitly contrastive (Schaffer, in preparation). What it takes, for instance, to know that there is a goldfinch in the garden differs if the contrasts are (i) goldfinch versus blue jay, (ii) goldfinch versus canary, or (iii) in the garden versus at the neighbor's. So compare:

*KQ1*: I know whether there is a goldfinch or a blue jay in the garden.

*KQ2*: I know whether there is a goldfinch or a canary in the garden.

*KQ3*: I know whether there is a goldfinch in the garden or at the neighbor's.

To know whether there is a goldfinch or a blue jay in the garden, a glance at the bird will suffice. Virtually anyone can answer that question. To know whether there is a goldfinch or a canary in the garden, some expertise with birds is required. Perhaps only the expert birder can answer that question. While to know whether there is a goldfinch in the garden or at the neighbor's, one must know the lay of the land. Perhaps only the homeowner can answer that question. Overall, what it takes to know that there is a goldfinch in the garden varies, depending on the question.

Just as what it takes to know that there is a goldfinch in the garden varies depending on the question, so what it takes to know whether one has hands varies. In this light, here is the contextualist treatment of external world skepticism in the language of knowledge-*wh*:

*KQ4*: I know whether I have hands or stumps.

*KQ5*: I know whether I have hands or am just a brain-in-a-vat.



The contextualist line is that *KQ4* is true and *KQ5* false. To know whether I have hands or stumps a glance at my upper extremities will suffice. While to know whether I have hands or am just a brain-in-a-vat, perhaps nothing will suffice.

There is nothing unusual about the verb ‘knows’ here. Many other propositional attitude verbs exhibit the same question (or contrast) sensitivity. Thus compare:

*P1*: I prefer that Bush is president [rather than Kerry].

*P2*: I prefer that Bush is president [rather than Cheney].

For a speaker such as myself whose overall preference ranking is Kerry, then Bush, then Cheney, an utterance of *P1* would be false, but *P2* true. Or compare:

*R1*: I regret who is president [as between Bush and Kerry].

*R2*: I regret who is president [as between Bush and Cheney].

For a speaker such as myself, an utterance of *R1* would be true, but *R2* false. Or suppose that I remember well that Kerry lost, but can’t quite recall which of those Republicans on the ticket won. Then compare:

*F1*: I forgot whether Bush is president [or Kerry].

*F2*: I forgot whether Bush is president [or Cheney].

In the scenario as described, an utterance of *F1* would be false, since I do remember that Kerry lost; but an utterance of *F2* would be true, since I can’t recall which of the Republicans won.

The cases of ‘regrets’ and ‘forgets’ are especially pertinent, since they are of the same lexical kind as ‘knows’ – all are factive attitude verbs that take both *that*- and *wh*-headed complements. It would be surprising if ‘knows’ behaved differently from its lexical kin.

So in conclusion, the strongest case for contextualism – or at least, for a theory in which the knowledge relation is relative to a set of contrasts – is based on the sensitivity of our intuitions to what is in question for the attributor. The subject is irrelevant. No form of invariantism can capture this

data. There is no substitute for the contextualist semantics – the contextualist data requires contextualism.<sup>16</sup>

#### NOTES

<sup>1</sup> Hawthorne (2004) argues that SSI is the best fit for the role of knowledge in assertion and practical reasoning. I address the role of knowledge in assertion in Section 3.

<sup>2</sup> The differences in both what is at stake for the subject, and what alternatives are made salient to us, are found in all extant presentations of these cases. Brown (forthcoming) notes this problem as well. Thus Fantl and McGrath insert the following additional material into their high-stakes *Train Case 2*: Maybe the ticket-seller misunderstood his question. Maybe he misunderstood the answer. Who knows when he bought the ticket? (2002, p. 68).

<sup>3</sup> Actually this is the prediction of the Hawthorne and Stanley versions of SSI. The Fantl and McGrath version predicts differently, since on their account what is crucial is not the stakes *per se*, but rather the stakes relative to the costs of further inquiry. Since *Low and Slow* and *High and Fast* differ over the costs of further inquiry, the Fantl and McGrath version may fit the intuitions as I report them for this case.

<sup>4</sup> Here and throughout I am using ‘contextualism’ in a broad sense, to cover not just standard contextualist approaches on which ‘knows’ is an indexical (Cohen, 1988, p. 97; DeRose, 1992, p. 920; Lewis, 1996), but also contrastivist approaches on which ‘knows’ invariantly denotes a ternary relation, whose third argument place (for contrasts) is set by context when left implicit (Morton and Karjalainen, 2003; Blaauw, 2004; Sinnott-Armstrong, 2004; Schaffer 2004, 2005a). See Section 4 for further discussion.

<sup>5</sup> This second objection also applies to the insensitive invariantist explanation of stakes-shifting provided by Bach (2005). Bach suggests that *High-on-Low* involves a performance error, in which the attributor panics, demanding more evidence than knowledge requires. This objection shows that the intuitions arise even when the attributor is completely unpanicked.

<sup>6</sup> There are further problems with the SSI semantics. For instance, there are the *embarrassing counterfactuals*. Claims such as the following come out true: ‘if only I had less at stake, I would have known’ and ‘if the stakes had been higher, I would not have known’ (Hawthorne, 2004, p. 166). There is also the cross-classification of skeptical scenarios. Brain in a vat scenarios like *The Matrix*, in which your loved ones are killed or enslaved turn out more epistemically difficult than ‘friendly’ BIV scenarios in which your loved ones are nurtured. I will not pursue these complaints

further here, but they ought to figure into one's overall assessment of the SSI semantics.

<sup>7</sup> This social role of the expert is connected to Craig's (1990) conception of the role of knowledge attribution, as identifying reliable informants.

<sup>8</sup> Along these lines, we also field questions with 'don't ask me – I don't know!' or 'ask her, she knows.' As Jon Kvanvig (personal communication) notes, we sometimes ask direct questions ('What time is it?') and sometimes ask indirect-epistemic questions ('Do you know what time it is?') It seems that we use direct questions when we think that the subject knows, and indirect-epistemic questions otherwise. For instance, if I see you looking at your watch I am likely to ask directly 'what time is it?' while otherwise I am more likely to pose the question indirectly 'do you know what time it is?' On the presupposition that you will at least know whether you know, the 'only ask those who know' rule explains the data.

<sup>9</sup> In this vein, Reynolds recalls: "a truck driver once told me, rather emphatically, that a service station attendant should give directions to a place only if he *knows* where it is . . . . We expect that people who don't know won't try to answer" (2002, p. 140).

<sup>10</sup> For sophisticated developments of this idea, see Higginbotham (1996) and Groenendijk and Stokhof (1997).

<sup>11</sup> Judging by what contestants on such shows say, the stakes seem to matter only when the subject does not know the answer, but just has some inkling. It is in these cases that she will factor the stakes (together with her level of confidence) into her decision. So it seems that what is at stake for the subject is not relevant to knowledge, but only relevant to decision making in the presence of doubt.

<sup>12</sup> Perhaps the best reply for the SSI-er would be to deny that knowledge is itself valuable. Kvanvig (2003) argues that knowledge itself adds no value to factors such as true belief, subjective justification, and cognitive virtue (in particular, he argues that whatever further conditions are needed to avoid Gettier problems will not be value-adding). So the SSI-er might follow Kvanvig's lead, and maintain that the stakes-sensitive component of knowledge is not value-adding.

<sup>13</sup> See Schaffer (2005b) for further discussion of what parameter shifts with context.

<sup>14</sup> Thus see, for instance, the contrastive theory of focus developed in Rooth (1992).

<sup>15</sup> Dretske may have been the first to note the impact of contrast shifts on knowledge ascriptions: "someone claiming to know that Clyde *sold* his typewriter to Alex is not (necessarily) claiming the same thing as one who claims to know that Clyde *sold* his typewriter *to Alex*. . . . A person who knows that Clyde *sold* his typewriter to Alex must be able to rule out the possibility that he *gave* it to him, or that he *loaned* it to him. . . . But he needs only a nominal justification, if he needs any justification at all, for

thinking it was Alex to whom he sold it" (1981, p. 373).

<sup>16</sup> Thanks to Kent Bach, Anthony Brueckner, John Hawthorne, Jim Joyce, Jon Kvanvig, John MacFarlane, Ram Neta, Walter Sinnott-Armstrong, Jason Stanley, and the Pacific APA audience.

#### REFERENCES

- Bach, K. (2005): 'The Emperor's New Knows', in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, Oxford: Oxford University Press.
- Blaauw, M. (2004): '*Contrastivism: Reconciling Skeptical Doubt with Ordinary Knowledge*', Amsterdam, Dissertation Free University of Amsterdam.
- Brown, J.: 'Contextualism and Warranted Assertibility Manoeuvres', *Philosophical Studies*. (forthcoming).
- Castañeda, H.-N. (1980): 'The Theory of Questions, Epistemic Powers, and the Indexical Theory of Knowledge', in P. French, T. Uehling Jr. and H. Wettstein (eds.), *Midwest Studies in Philosophy V: Studies in Epistemology* (pp. 193–238), Minneapolis: University of Minnesota Press.
- Cohen, S. (1988): 'How to be a Fallibilist', *Philosophical Perspectives* 2, 91–123.
- Cohen, S. (1999): 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives* 13, 57–89.
- Craig, E. (1990): *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*, Oxford: Oxford University Press.
- DeRose, K. (1992): 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research* 52, 913–929.
- DeRose, K. (1995): 'Solving the Skeptical Problem', *Philosophical Review* 104, 1–52.
- De Rose, K. (2004): 'The Problem with Subject-Sensitive Invariantism', *Philosophy and Phenomenological Research* 68, 346–350.
- De Rose, K.: 'The Ordinary Language Basis for Contextualism and the New Invariantism', *Philosophical Quarterly* (forthcoming).
- Dretske, F. (1981): 'The Pragmatic Dimension of Knowledge', *Philosophical Studies* 40, 363–378.
- Fantl, J. and McGrath, M. (2002): 'Evidence, Pragmatics, and Justification', *Philosophical Review* 111, 67–94.
- Groenendijk, J. and Stokhof, M. (1997): 'Questions', in J.F.A.K. Benthem van and A.G.B. ter Meulen (eds.), *Handbook of Logic and Language* (pp. 1055–1124), Amsterdam: Elsevier.
- Hawthorne, J. (2004): *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Higginbotham, J. (1996): 'The Semantics of Questions', in S. Lappin (ed.),

- The Handbook of Contemporary Semantic Theory* (pp. 361–383), Oxford: Oxford University Press.
- Hookway, C. (1996): ‘Questions of Context’, *Proceedings of the Aristotelian Society* 96, 1–16.
- Kvanvig, J. (2003): *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: Cambridge University Press.
- Lewis, D. (1996): ‘Elusive Knowledge’, *Australian Journal of Philosophy* 74, 549–567.
- MacFarlane, J. (2005): ‘Knowledge Laundering: Testimony and Sensitive Invariantism’, *Analysis* 65, 132–138.
- Morton, A. and Karjalainen, A. (2003): ‘Contrastive Knowledge’, *Philosophical Explorations* 6, 74–89.
- Reynolds, S. (2002): ‘Testimony, Knowledge, and Epistemic Goals’, *Philosophical Studies* 110, 139–161.
- Rooth, M. (1992): ‘A Theory of Focus Interpretation’, *Natural Language Semantics* 1, 75–116.
- Schaffer, J. (2004): ‘From Contextualism to Contrastivism’, *Philosophical Studies* 119, 73–103.
- Schaffer, J. (2005a): ‘Contrastive Knowledge’, *Oxford Studies in Epistemology* 1.
- Schaffer, J. (2005b): “What Shifts? Thresholds, Standards, or Alternatives?”, in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, Oxford: Oxford University Press.
- Schaffer, J.: “Knowing the Answer” (in preparation).
- Sinnott-Armstrong, W. (2004): *Pyrrhonian Skepticism*, Oxford: Oxford University Press.
- Stalnaker, R. (1999a): ‘Assertion’, in *Context and Content* (pp. 78–95), Oxford: Oxford University Press.
- Stalnaker, R. (1999b): ‘On the Representation of Context’, in *Context and Content* (pp. 96–113), Oxford: Oxford University Press.
- Stanley, J.: *Knowledge and Interest*, Oxford: Oxford UP (forthcoming).
- Williamson, T. (2000): *Knowledge and its Limits*, Oxford: Oxford University Press.
- Williamson, T.: ‘Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge’, *Philosophical Quarterly*. (forthcoming).

*Department of Philosophy*  
*University of Massachusetts-Amherst*  
*Amherst, MA 01003*  
*E-mail: schaffer@philos.umass.edu*