

KNOWLEDGE IN THE IMAGE OF ASSERTION

Jonathan Schaffer
ANU & Arché

How must knowledge be formed, if made in the image of assertion? That is, given that knowledge plays the normative role of governing what one may assert, what can be inferred about the structure of the knowledge relation from this role? I will argue that what one may assert is sensitive to the question under discussion, and conclude that what one knows must be relative to a question. In short, knowledge in the image of assertion is question-relative knowledge.

1. Assertion

How must knowledge be formed, if made in the image of assertion? It will prove useful to begin with a discussion of what assertion *does*. For I am after the role that knowledge plays in the norms of assertion, and these norms are licenses to do what assertion does. I will argue that all contexts include a *question under discussion*, and what an assertion does is *answer the question*.

1.1 Assertion and Contrast: Stalnaker

What is it that assertion does? According to the orthodox Stalnakerian account, what assertion does is *aim to reduce the context set*. The context set is, in Stalnaker's words "the set of possible worlds recognized by the speaker to be the 'live options' relevant to the conversation" (1999a: 84–5). This may be represented by a set of possible worlds "which include all the situations among which speakers intend to distinguish with their speech acts" (1999b: 99). For instance, if we are presupposing that exactly one bird is in

2 Jonathan Schaffer

the garden, but wondering whether it is a goldfinch, raven, or canary, then the context set would be the set of worlds S containing the goldfinch-in-the-garden, raven-in-the-garden, and canary-in-the-garden worlds.

The assertion has a content p , which may also be represented as a set of possible worlds, namely those in which p is true. So if I assert “It is a goldfinch,” the content may be understood in terms of the set of worlds in which it (the bird in question) is a goldfinch.

The way in which assertion aims to reduce the context set then goes as follows. Take the intersection of the context set and the content set. Call that *the update*. What an assertion aims to accomplish in a given context is to reduce the context set to the update:

To make an assertion is to reduce the context set in a particular way. . . The particular way in which the context set is reduced is that all of the possible situations incompatible with what is said are eliminated. . . [T]he essential effect of an assertion is to change the presuppositions of the participants in the conversation by adding the content of what is asserted to what is presupposed. (Stalnaker 1999a: 86)

For example, if I assert “It is a goldfinch” in a context aptly represented by S , my assertion aims to eliminate the blue jay, raven, and canary options, and reduce S down to the set of worlds in the intersection $S\uparrow$, which only contains goldfinch-in-the-garden worlds.


To illustrate, think of the content p as a set of worlds P :




Think of the context S as another set of worlds:



Think of the update $S\uparrow$ as the intersection of the content and context sets ($S\uparrow = P \cap S$):



One last bit of terminology: call the remainder $S\downarrow$ of the context set *the contrast*. The contrast is the set of $\sim p$ -worlds in S ($S\downarrow = S - P$):



What an assertion of p in a context aptly represented by S aims to do is to eliminate the contrast $S\downarrow$, and thereby reduce S to $S\uparrow$.

So on the Stalnakerian model, what an assertion does is aim to eliminate the contrast, and thus reduce the context to the update. License to assert is license to do what assertion does. So on the Stalnakerian model, license to assert is license to aim at eliminating the contrast.

1.2 Assertion and Question Under Discussion: Carlson, Ginzburg, and Roberts

According to Carlson (1983), Ginzburg (1996), and Roberts (2004), contexts include a *question under discussion*. The question under discussion is posited as one of the entries on the ‘conversational scoreboard’ (Lewis 1979), reflecting what is being addressed at that point in the conversation. More fully, the scoreboard might include such entries as: (i) a stack of questions (the topmost of which is the question under discussion), (ii) a set of ‘file cards’ bearing information on discourse referents, and (iii) a set of sets of goals (one set per interlocutor), *inter alia* (c.f. Roberts 2004: 215).

The question under discussion is posited to explain phenomena such as distant ellipsis, focus placement, and felicitous topicalization. For instance, Ginzburg cites an actual discussion which starts with the question of how old someone is, and ends many turns later with the elliptical utterance “seventy two.” The question must be available to interpret the ellipsis. As Ginzburg explains:


[S]ince the discussion of a single question can last over several turns, and elliptical contributions are possible, in principle, arbitrarily far away from the turn in which the question was posed, what will be needed is a notion of context which can express the fact that a particular question is (still) under discussion, and hence its associated relation is (still) salient. (1996: 415)

The question under discussion is also used to explain and clarify the Gricean maxims, rules of turn taking, and discourse coherence. Thus Carlson (1983; c.f. Roberts 2004: 209) speaks of the discourse game as built around questions (setup moves) and answers (payoff moves), and Roberts explains and clarifies the Gricean maxim of *relevance* in terms of addressing the question:


A move m is RELEVANT to the question under discussion q iff m either introduces a partial answer to q (m is an assertion) or is part of a strategy to answer q (m is a question subordinate to q or an imperative whose realization would plausibly help to answer q). (2004: 216)

In short, to be relevant is *to speak to the question*.

Models of context that posit a question under discussion may be seen as extending the Stalnakerian model, by adding internal partitional structure to context sets. According to the orthodox view of questions (c.f. Higginbotham 1996, Groenendijk and Stokhof 1997), questions are *partitioned regions of logical space*, where each cell of the partition corresponds to a possible answer to the question. For instance, the question “What bird is in the garden?”, asked in a context in which goldfinch, raven, and canary are in the domain, and in which it is presupposed that exactly one bird is in the garden, would denote the partition $\Pi = \{S_1 = \{w: \text{there is a goldfinch (and no other bird) in the garden}\}, S_2 = \{w: \text{there is a raven (and no other bird) in the garden}\}, S_3 = \{w: \text{there is a canary (and no other bird) in the garden}\}\}$. Thus instead of one big context set S as on the Stalnakerian model (§1.1):




The question under discussion adds internal partitional structure to the representation of context:




So understood, the question under discussion model is perhaps truer to Stalnaker’s own conception of context as “the live options” and “the situations among which speakers intend to distinguish”. The speakers might well intend to distinguish between goldfinch and raven worlds, but not within goldfinch worlds. The partitional structure marks this intention.

The content p is still to be modeled as a set of worlds. The update can still be modeled by intersection of p with the worlds in the context partition Π . Where the update is to a single cell of the original partition, call this a *complete answer*. For instance, if one asserts “It is a goldfinch” in a context aptly modeled by Π , one selects cell S_1 as the complete answer to the question under discussion (using shading to represent what the content eliminates):



Not all answers are complete answers—some are *partial*, leaving multiple cells uneliminated, such as “It is not a raven”:



Others are *rejective*, breaking the context partition by eliminating every cell, such as “It is a squirrel.”¹

In what follows I will work with the question under discussion model. What an assertion of p in a context aptly represented by question Π aims to do then is to answer the question. As Roberts says, assertions provide “the answers to questions” (2004: 209). So instead of saying that an assertion aims at eliminating the contrast (as on the Stalnaker model), one may say:

AA: An assertion aims to answer the question under discussion

License to assert is license to do what assertion does. So license to assert is license to aim at answering the question.

1.3 Assertibility is Sensitive to the Question

If license to assert is license to aim at answering the question under discussion, one might wonder if license to assert might depend on what that question is. I will now argue that license to assert does so depend. Thus consider the following example, borrowed from Austin (via DeRose 2002: 169):

Find Fred: We are looking for Fred. We find Fred’s hat hanging by the door, and on that basis I say that Fred is here.

Do I have license to make this assertion? *Maybe*. In contexts where we are presupposing that Fred’s hat indicates his presence, *yes* (all else equal). Such an assertion would be perfectly in order. But in other contexts *no*—imagine, for instance, that we are wondering if Fred might have left without his hat. Then the mere sighting of Fred’s hat hanging by the door will not position me to say whether Fred is here or not. Thus whether or not I may assert that Fred is here depends on what is in question. It depends in particular on whether the question includes options in which Fred and his hat have parted ways.

Here is a second example of license to assert depending on the question, inspired by Dretske (1970):

Stop Thief: Black has stolen the opals from the locked safe. The detective arrives on the crime scene, finds Black's fingerprints all over the safe, and on that basis reports that Black has stolen the opals.

Did the detective assert appropriately? *It depends on the question.* If the question concerned who stole the opals, then the answer is *yes* (all else equal)—the fingerprints reveal that it was Black. But if the question concerned what Black stole then the answer is *no*—the fingerprints on the safe do not show what was removed from the safe. The fingerprints reveal that it was Black who did the deed, but they do not reveal what deed Black did.

Here is a third example of question-relativity, also drawn from Austin (1946):

The Goldfinch in the Garden: There is a goldfinch in the garden. Ann and Ben happen to be strolling by. Ann spots the golden plumage, and thus remarks that there is a goldfinch in the garden.

Did Ann assert properly? If the question is whether there is a goldfinch in the garden or a raven, then *yes* (all else equal)—Ann has spotted the golden plumage and so can rule out the prospect of raven. But if the question is whether there is a goldfinch in the garden or a canary, then *no*—not being an expert birder, Ann cannot tell a goldfinch apart from a canary. Whereas if the question is whether there is a goldfinch in the garden or at the neighbor's, then *no* as well—perhaps only the homeowner is aware of where exactly the property lines lie. Thus whether Ann may assert that there is a goldfinch in the garden depends on the question.

For a fourth and final example, consider the following easy question:

Who was the thirteenth president of the United States?
(a) Millard Fillmore
(b) Hillary Clinton


Your average student, faced with this question, can appropriately assert that Millard Fillmore was the thirteenth president of the United States. But consider the following hard question:

Who was the thirteenth president of the United States?
(a) Millard Fillmore
(b) Zachary Taylor


Your average student, faced with the hard question, cannot appropriately assert that Millard Fillmore was the thirteenth president of the United

States. As anyone who has made a multiple-choice exam will recognize, the alternatives matter. Ask an easy question and your students will sing out the answer; but add some devious alternatives and your students will sheepishly stay silent, or recklessly guess.²


These cases have a common structure. Graphically, there is the content asserted (modeled as the set P), and various possible options (Q_1 and Q_2):




For instance, in *Find Fred*, $P = \{w: \text{Fred is here with his hat in } w\}$, $Q_1 = \{w: \text{Fred is elsewhere with his hat in } w\}$, and $Q_2 = \{w: \text{Fred is elsewhere without his hat in } w\}$. Now the speaker has evidence sufficient to rule out one option but not the other. Use shading to represent what the speaker's evidence rules out, this may be depicted as:



For instance, the sight of Fred's hat here rules out Q_1 but not Q_2 . Now in a context aptly modeled by the partition $\Pi_{easy} = \{P, Q_1\}$, the assertion of p aims to answer the easy question, which is as the evidence warrants:



But in a context aptly modeled by $\Pi_{hard} = \{P, Q_2\}$, the assertion of p aims to answer the hard question, which the evidence does not warrant:



And so in these different contexts, the assertion that Fred is here aims to answer different questions, and the evidence validates the former aim but not the latter.

Putting this together: to have license to assert is to have license to aim at answering the question. Whether one has license to aim at answering the question depends on what the question is, because one might be in position to answer some questions but not others. Thus I submit:

AQ: Whether one may assert that p in context c depends on the question under discussion in c .

When one can answer the question, the assertion is licensed (all else equal). But when one cannot answer the question, the assertion is not permitted.

2. Knowledge

So far I have focused on assertion, and argued that what one may assert is sensitive to the question under discussion (as per AQ). I am building towards the claim that knowledge is relative to a question, so I now turn to links between what one may assert and what one knows.

2.1 Knowledge as a Norm of Assertion: Unger and Williamson

How are assertion and knowledge linked? On the Williamsonian view (the spirit of which I will uphold), knowledge plays a normative role governing what one may assert. The norm in question may be glossed as: *say it only if you know it*; or, *know whereof you speak*.³

The justification for treating knowledge as a norm of assertion emerges from the following four points, culled from Unger (1975: 250–71) and Williamson (2000: 238–69):

Critique: One who asserts that p but does not know that p is liable to *criticism*.

For instance, even the Gettier-ized subject who glances at a clock stopped 24 hours ago and asserts: “It is 10am” is open to censure. He does not know what he is talking about.

Defect: Assertions of the form: “ p , and I do not know that p ” are *defective*.

For instance, if I assert “Grass is green, and I do not know that grass is green,” then I have asserted defectively, even though the proposition expressed may well be true. The defectiveness involved is comparable to that found in the Moore-paradoxical “Grass is green, and I do not believe that grass is green.”

Challenge: An assertion that p may be *challenged* by “How do you know” questions.

As Austin (1946) notes, one who asserts: “There is a goldfinch in the garden” is thereby open to challenges, such as: “How do you know that it is not a canary?” and “How do you know that it is not at the neighbor’s?”

Authority: Asserting, like commanding, requires *authority*.

Just as commanding requires social authority, so asserting requires epistemic authority. The ignorant asserter is as out of line as the student who dismisses class.

These four points of linkage between knowledge and assertion cry out for explanation. Williamson (2000: 243) explains these points via the following knowledge norm:

KA: *S* ought: assert that *p* only if *s* knows that *p*.

Given KA, *Critique* is explained via the fact that ignorant asserter has not done as she ought. *Defect* is explained from the facts that (i) one cannot have license to assert a conjunction without license to assert each conjunct, and (ii) the license to assert the *p*-conjunct is that one knows that *p*, where (iii) this license contradicts the content of the “I do not know that *p*” conjunct. *Challenge* is explained by the fact that “How do you know” questions challenge one’s license to assert. *Authority* is explained by knowledge providing epistemic authority. So far, so good.^{4,5}

But I think that KA suffers from two (interrelated) problems. First, KA is *blind to the act*. The norm it encodes for assertion only concerns the content of what is asserted (*p*), ignoring the effects of the act of asserting that *p* in a given context. One might have expected the norms for speech acts to be norms for action, that concern not merely the content asserted, but also what the speakers *does* by asserting such a content in a given context (which is to aim at answering the question: §1.2). The norms for assertion should be licenses to do what assertion does.

My second concern with KA is that it *fails to explain AQ*. The only licenses KA issues turn on whether *s* knows that *p*. This license is not connected to the question. So for instance consider *The Goldfinch in the Garden*, and ask whether Ann knows that there is a goldfinch in the garden, or not. If *yes*, then KA does not explain why Ann cannot assert that there is a goldfinch in the garden, when considering the (hard) canary question. If *no*, then KA entails that Ann ought not assert that there is a goldfinch in the garden, even when considering the (easy) raven question.

Note that it is compatible with KA that (i) in all the question-sensitivity cases, the asserter indeed knows whereof she speaks; but (ii) the question plays a role in some *other norm* governing assertion, which other norm then explains AQ. (I am only saying that KA itself fails to explain AQ, I am *not* saying that KA precludes the existence of any *other* explanation for AQ.) But clearly a formulation of the knowledge norm that could itself explain AQ would be preferable. First, it would lift the theoretical burden of having to supply a further norm that does explain AQ. Second, it seems to me that the question-sensitivity of what one may assert is *keyed to what one knows*.

Think of a student facing the easy question of whether Millard Fillmore or Hillary Clinton was the thirteenth US president. The reason why the student is licensed to assert is that *she knows the answer*. Now think of a different student facing the hard question of whether Millard Fillmore or Zachary Taylor was the thirteenth US president. The reason why this student is not licensed to assert is that *he does not know the answer*.⁶ So I would suggest that the explanation for AQ should if possible proceed *via the knowledge norm*.

2.2 The Question-Relative Knowledge Norm

So how must knowledge be formed, if made in the image of assertion? If what one may assert depends on the question (§1.3), and what one may assert is linked to what one knows (§2.1), this suggests that what one knows should be treated as relative to a question. Knowledge is itself a question-relative notion. I am now in position to defend the following account of the knowledge norm on assertion, which embeds a question-relative notion of knowledge:


KQ: *S* ought: assert that *p* in context *c* only if *s* knows the answer (*p*) to the question under discussion in *c*.

In other words: license to aim at answering the question (which is what assertion does) requires knowing the answer to the question. That is: *offer an answer to the question only if you know the answer to the question*. In short: *say the answer only if you know the answer*.


The epistemic state embedded in KQ is a *ternary* relation between a subject, a proposition, and a question.⁷ It requires ruling out the non-*p* options: if *s* knows the answer to the question, then *s* must be in position to rule out all but the true answer. This allows *s* to know the answer to one question but not another, even where *p* is the true answer to both. Such is the state of the typical student who knows the answer to the question of whether Millard Fillmore or Hillary Clinton was the thirteenth US president, but does not know the answer to the question of whether Millard Fillmore or Zachary Taylor was the thirteenth US president. Such a student can rule out the Clinton option but not the Taylor option. Likewise this is the state of the detective in *Stop Thief*, who knows the answer to the question of who stole the rubies, but does not know the answer to the question of what Black stole. The detective can use the fingerprint evidence to rule out alternative thieves, but not to rule out alternative loot.

The main motivation for KQ is that *KQ fits what assertion does*. License to assert is license to do what assertion does, and what assertion does is to aim at answering the question. Hence the knowledge required for assertion

is knowledge of the answer to the question. Graphically, in a context in which the question under discussion is $\Pi = \{\{w: w \text{ is a } p\text{-world}\}, \{w: w \text{ is a } q\text{-world}\}\}$, the aim of asserting p is to answer Π :



And the evidential requirement on $K_{sp}\Pi$ is the ruling out of the shaded portion:



Question-relative knowledge is thus knowledge in the image of assertion.

Would anything be lost, in moving from KA to KQ? I think not; or at least, I would maintain that KQ preserves all the explanatory virtues of KA, with respect to *Critique*, *Defect*, *Challenge*, and *Authority* (§2.2). As to *Critique*, the critique of the ignorant asserter is that he does not know the answer to the question. The Gettierized subject staring at the broken clock, for instance, does not know what time it is. He is in no position to rule out time alternatives. The typical student does not know whether Millard Fillmore or Zachary Taylor was the thirteenth US president. She can only guess.

As to *Defect*, what is defective about “ p and I do not know that p ” assertions is that the knowledge required to assert the first conjunct is labeled a guess in the second. One cannot have license to assert a conjunction without license to assert each conjunct. The license to assert the p -conjunct in the context of utterance c , is that one knows the answer (p) to the question under discussion in c . But the content of the “I do not know that p ” conjunct is that one does not know that p relative to the question under discussion. Absent some reason to except micro-contextual shift, this will be the same question. And so the license to assert the first conjunct is incompatible with the content of the second conjunct.

Turning to *Challenge*, “How do you know” questions challenge one’s license to answer the question. They invoke further options, and challenge whether these have been properly eliminated. Indeed, KQ might provide a better explanation of *Challenge* than KA did. For some “How do you know” queries are more challenging than others. Imagine a witness testifying in court that she saw Sam flee the scene of the crime. The defense attorney may well challenge the witness with “How do you know that it was Sam rather than his brother Tom?” But the witness cannot be challenged with “How do you know that you were not a brain-in-a-vat hallucinating the whole scene?” That

option is not under discussion in the courtroom. The lawyer who tried such a challenge would be ignored. KQ explains *which* “How do you know” queries are challenging—they are the ones that concern *the options in question*.

Authority, finally, is explained by the epistemic authority that knowledge provides. Indeed, KQ might provide a better explanation of *Authority* than KA. Since what an assertion does is to aim at answering the question, the authority required in KQ is more fitting to the act. I thus conclude that KQ is at least as good as KA (if not better) in connecting what one knows to what one may assert.

Moreover, KQ avoids the two criticisms of KA given in §2.1. First, KQ is alive to the act of assertion—indeed, KQ is built around the leading linguistic account of what assertion does (§1.2).

Second, KQ explains AQ. The reason why assertibility is sensitive to the question is because (i) assertibility is governed by a knowledge norm, and (ii) knowledge is relative to a question. For instance, in *Find Fred*, if the question is whether Fred is elsewhere with his hat, then seeing his hat hanging by the door puts me in position to eliminate the false option, and thereby provides me the basis for the knowledge required to assert that Fred is here. But if the question is whether Fred is elsewhere having left his hat behind, then seeing his hat hanging by the door puts me in no position to eliminate the false option, and thereby does not provide me any basis for the knowledge required to assert that Fred is here.

Of course KQ might be held to be implausible, insofar as it embeds a question-relative knowledge relation, and insofar as that is held to be implausible. I admit that question-relativity is a radical thesis. I admit that it might seem implausible *if* one only thinks in terms of simple knowledge-*that* claims. So if one only considers claims like “I know that I have hands,” one might wonder *where is the question?* (To which the answer is: *on the conversational scoreboard, as the question under discussion*.) But one should also consider knowledge-*wh* claims, such as “I know whether Millard Fillmore or Zachary Taylor was the thirteenth US president.” In such claims, the question is sitting right on the surface. Question-relativity is really no more radical than the following claim: to know whether Millard Fillmore or Hillary Clinton was the thirteenth US president, is to know the answer to the question of whether Millard Fillmore or Hillary Clinton was the thirteenth US president.

3. Contextualism

There is a standing debate in the literature as to whether treating knowledge as a norm of assertion establishes contextualism (DeRose 2002), or refutes it (Hawthorne 2004). I will conclude by considering the bearing of KQ on this debate. What emerges is a vindication of a specific but unusual form of contextualism.

3.1 Contextualism Vindicated: DeRose and Blackson

I begin with DeRose's claim that treating knowledge as a norm of assertion establishes contextualism. DeRose begins by noting that the epistemic position to assert is not constant (citing cases akin to *Find Fred*: §1.3). He then argues that, since the position to assert is contextually variable, knowledge must swing alongside it:

The knowledge account of assertion provides a powerful argument for contextualism: If the standards for when one is in a position to warrantably assert that P are the same as those that comprise a truth-condition for "I know that P", then if the former vary with context, so do the latter. In short: The knowledge account of assertion together with the context-sensitivity of assertibility yields contextualism about knowledge. (2002: 187)

While I will ultimately agree with DeRose that treating knowledge as a norm of assertion leads to contextualism, I think DeRose's argument is invalid as stated (as pointed out in Blackson 2004). From the fact that the standards for warranted assertibility are not constant, contextualism does not yet follow. What matters is whether the inconstancy is due to *the situation of the subject* or *the linguistic context of the attributor*. Thus consider subject sensitive invariantism (hereafter SSI), as defended by Hawthorne (2004: §4) and Stanley (2005). SSI is an invariantist theory. Yet because SSI connects knowledge to the practical interests of the subject, it can be combined with KA to predict that the standards for assertibility are not constant, but rather vary with the practical situation of the subject.⁸

What makes it especially difficult to decide whether the norms of assertion favor contextualism or SSI is that the one who needs to know (the subject) is the speaker (the attributor). So with the norms of assertion there is no prospect of using 'third person' cases to pry contextualism and SSI apart. Rather one must appeal to the exact way in which what one may assert shifts. One must see whether the shift profile of assertibility matches the shift profile contextualism predicts in shifting with the linguistic context, or whether it matches the shift profile SSI predicts in shifting with what is at stake (or perhaps some other profile entirely).

I have argued that assertibility shifts with the question, as per AQ. This is a matter of the speaker's linguistic context, as contextualism predicts. Indeed, this is not at all a matter of what is at stake (which is what SSI predicts). None of the motivating cases for AQ (§1.3) even mentioned what anyone had at stake, nor did this need to be mentioned. In *Goldfinch in the Garden*, for instance, Ann presumably had the same very low level of (objective and subjective) interest in any of the various possible options. She was merely making conversation. So AQ, by keying assertibility to the speaker's *linguistic context*, provides the additional premise needed to vindicate the argument to contextualism.⁹

Though note that KQ calls for an unusual version of contextualism, on which knowledge is a ternary, question-relative relation $K_{sp}\Pi$. This differs from the usual formulations of contextualism in at least three main respects. First, what shifts is the value of a third argument (the question). This allows “knows” itself to remain *invariant* in semantic value, always expressing the one and only K relation—all that shifts is the value of Π . Second, the value of this third argument is fixed by a general and independently needed contextual parameter (the question under discussion). This means that no special rules of relevance for knowledge ascriptions need be invented. Third, what fixes the value of this third argument is the attributor’s linguistic context. This means that no subject factors play a role in determining the alternatives. As I have argued elsewhere, all of this represents an improvement over standard forms of contextualism, in which the semantic value of “knows” varies according to special rules of relevance that involve both attributor and subject factors.¹⁰ So while I take DeRose’s argument from the knowledge norm on assertion to contextualism to be vindicated, I do not take it to vindicate DeRose’s version of contextualism, nor any of the standard versions of contextualism now in the literature.

3.2 *Contextualism Defended: Hawthorne*

I conclude by considering an objection to a contextualist treatment of the norms of assertion, due to Hawthorne (2004). I will argue that this objection (if I have understood it properly) is not very objectionable—indeed the contextualist treatment is getting things right.¹¹

Hawthorne’s objection (as I understand it) stems from the fact that KQ is a *speaker-centered* norm, while contextualism is an *ascriber-dependent* theory. KQ ties what the speaker may assert into the question under discussion in the speaker’s context. But contextualism ties the question relevant to knowledge ascription into the ascriber’s context. So given KQ, when speaker and ascriber are discussing different questions, the powers to answer that warrant assertion may diverge from those that ground knowledge ascription. The ascriber (suitably informed of the contextualist semantics) may even recognize such divergence, and thus assert oddities such as: “People often flat-out assert things that they do not know to be true but are not thereby subject to criticism” (Hawthorne 2004: 87), and “There are things people know but ought not to assert because their epistemic position is not strong enough with respect to those things” (Hawthorne 2004: 87). Indeed these assertions sound on surface as if they were (respectively) denials of KA, and of a plausible converse thesis to the effect that knowledge *suffices* for being in an epistemic position to assert. Thus Hawthorne charges that, under contextualism, “the link between knowledge and assertibility has been severed” (2004: 87).

Two variants of *Stop Thief* may be used to illustrate this alleged severing. In the one direction, consider:

Good Cop: Black has stolen opals from the locked safe. Two different detectives, with very different presuppositions about what has happened, have been dispatched to the scene of the crime. Detective Hu has been asked to determine whether Black or Red has stolen opals. Hu finds Black's fingerprints all over the safe, and on that basis reports that Black has stolen opals. Detective Watt has been asked to determine whether Black has stolen opals or rubies. Watt watches Hu in action, and announces: "Detective Hu is in a perfectly good epistemic position to assert that Black has stolen opals, but Hu does not know that Black has stolen opals."

Detective Watt's concluding announcement sounds odd. Indeed it sounds like an explicit denial of knowledge as a norm of assertion ('Hu may say it but he does not know it'). Yet given KQ and the fact that the option Hu needed to eliminate was eliminated via the fingerprint evidence, Hu is epistemically positioned to assert that Black has stolen opals. The first conjunct of Watt's announcement comes out true. And given contextualism and the fact that Watt is concerned with the ruby option that Hu has not eliminated, Watt should deny that Hu knows that Black has stolen opals. The second conjunct comes out true as well.

In the other direction, consider:

Bad Cop: Black has stolen opals from the locked safe. Two different detectives, with very different presuppositions about what has happened, have been dispatched to the scene of the crime. Detective Hu has been asked to determine whether Black or Red has stolen opals. [So far matters are just as they were in *Good Cop*.] This time Hu finds Black's fingerprints all over the safe, but does not yet assert anything. Now Detective Watt (as in *Good Cop*) has been asked to determine whether Black has stolen opals or rubies. This time Watt watches Hu in action, and asserts that Black has stolen opals. Hu hears this and announces: "Detective Watt knows that Black has stolen opals, but he is in no epistemic position to assert that Black has stolen opals."

Detective Hu's concluding announcement sounds odd. Indeed it sounds like an explicit denial of a plausible converse of KA, to the effect that knowledge *suffices* for epistemic position to assert ('Watt knows it but he is in no position to say it'). Yet given KQ and the fact that the option Watt needed to eliminate was not eliminated via the fingerprint evidence, Watt is in no epistemic position to assert. The first conjunct of Hu's announcement comes out true. And given contextualism and the fact that Hu is concerned with the Red alternative that has been eliminated, Hu should affirm that Watt knows that Black has stolen opals. The second conjunct comes out true as well.

So runs the objection. By way of reply, I agree that the assertions that conclude *Good Cop* and *Bad Cop* are odd, but will argue that this is merely an artifact of their odd wording, and not a product of their falsity. To see that the wording is odd (independent of any views on truth or falsity), note that these assertions *require the implicit options to shift mid-sentence*. For instance, consider Watt's odd assertion: "Detective Hu is in a perfectly good epistemic position to assert that Black has stolen opals, but Hu does not know that Black has stolen opals." The first conjunct comes out true because Hu knows the answer to the question of who stole the opals (speaker's question), and the second conjunct comes out true because Hu does not know what Black stole (ascriber's question). In general, KQ allows that the powers to answer that warrant assertion may diverge from those that ground knowledge ascription, precisely because they may involve this sort of difference in questions under discussion.

Assertions that require implicit context shifts mid-sentence often feel odd and confusing, as can be shown via the contrastive notion of *preference*. Thus consider:

Just Desserts: Ann's preference ranking for ice cream is chocolate, then vanilla, then strawberry. She is offered a choice of either vanilla or strawberry. Ben comments: "It is rational for Ann to choose vanilla, though Ann does not prefer vanilla."

Ben's concluding comment ought to sound odd. Indeed, it seems that Ben is saying that it is rational to choose what one does not prefer (in seeming contradiction to the plausible principle *choose what you prefer*). But of course this is only a trick of wording. Ben's wording is odd and confusing because it requires the implicit contrast to shift mid-sentence (from *strawberry*, to *chocolate*). Far less confusing would have been for Ben to speak explicitly: "It is rational for Ann to choose vanilla *over strawberry*, though Ann does not prefer vanilla *to chocolate*."

Does anything in *Just Desserts* show that contrastive preference has been 'severed' from rational choice? *Of course not*. All that has been shown is that, if one is suitably devious, one can construct very odd sentences concerning contrastive preference and rational choice. This seems to me sufficient to explain away all the oddness of the concluding assertions of *Good Cop* and *Bad Cop*, and generally to defuse Hawthorne's objection (as I understand it).

Indeed, I would add that more explicit wording in the knowledge case would alleviate much of the felt oddness, just as it did in the preference case. For instance, when Watt concludes *Good Cop* with "Detective Hu is in a perfectly good epistemic position to assert that Black has stolen opals, but Hu does not know that Black has stolen opals," a more explicit rendition of the content expressed is that (i) Hu is in a perfectly good epistemic position to

answer the question of whether Black or Red has stolen the opals, but (ii) Hu does not know the answer to the question of whether Black has stolen opals or rubies. Both (i) and (ii) seem true. After all, Hu does indeed know whether Black or Red stole the opals, on the basis of finding Black's fingerprints. Further, Hu does not know what Black stole—he is in no position to answer the question of whether Black has stolen opals or rubies, since the fingerprints do not speak to this question. In short, Hu knows whether Black or Red has stolen opals, but does not know whether Black has stolen opals or rubies. (Recall Ann in *The Goldfinch in the Garden*, who knows whether there is a goldfinch or a raven in the garden, but does not know whether there is a goldfinch in the garden or at the neighbor's.) This is some positive evidence that the assertions in question (though confusingly phrased) were in fact *true*.

Pending further reason to think that anything has gone wrong in *Good Cop* and *Bad Cop* other than odd wording, and pending further objections, I conclude that the question-relative treatment is getting things right. What one may assert and what one knows both depend on the question. Knowledge in the image of assertion is question-relative knowledge.¹²

Notes

1. This is only a partial taxonomy of answers. For instance, there are also *restructuring answers* that divide partitions, such as "It is a lesser goldfinch." (These are perhaps best modeled as shifting the question under discussion to a more fine-grained one, to which they then serve as complete answers.)
2. In these cases it is crucial to distinguish (i) what the students can answer *before* hearing the question, from (ii) what the students can answer *after* hearing the question. Hearing the question from the teacher may provide additional information. (For instance, if one hears both the easy question and the hard question, one can then extract enough information to answer the hard question!) I am primarily interested in what the students can answer *before* hearing the question, given their current information.
3. In a similar vein, Unger hypothesizes: "If S asserts, states, or declares that p , then he not only represents it as being the case that p , but he represents it as being the case that he *knows* that p " (1975: 253). To assert what one does not know is thus to falsely represent oneself. That would be bad.
4. Note that KA itself is neutral with respect to some further Williamsonian theses, including (i) the *uniqueness* claim that KA is *the one and only* norm specifically applicable to assertion (2000: 241), and (ii) the *constitutivity* claim that assertion is *defined* as the speech act φ governed by the rule: S ought: φ that p only if s knows that p . (2000: 238).
5. Though see Weiner 2005 and Lackey 2007 for some alleged counterexamples to any knowledge norm of assertion (including not just KA, but also the norm I will propose in §2.2). I lack the space to contest Weiner's and Lackey's cases here. The skeptical reader is welcome to treat the argument of the main text as a

conditional argument: *if* knowledge is a norm of assertion, *then* knowledge must be a question-relative relation.

6. It is worth keeping in mind that we use multiple-choice questions precisely to test what our students know. (Think of the F student as knowing the answer to neither the easy nor hard question, the C student as knowing the answer to the easy question only, and the A student as knowing the answer to both.)
7. The question-relative account of knowledge is best understood as a form of the *relevant alternatives* theory of knowledge (introduced by Austin 1946), on which the relevant alternatives are the remaining possible answers to the question. Thus Dretske states: “To know that *x* is *A* is to know that *x* is *A* within a framework of relevant alternatives, *B*, *C*, and *D*. This set of contrasts. . . serve to define what it is that is known. . .” (1970: 1022) Versions of this theory—which I elsewhere call the *contrastive view of knowledge*—have since been defended by Johnsen (2001), Morton and Karjalainen (2003), Sinnott-Armstrong (2004), Schaffer (2004, 2006b, 2007), and Blaauw (2004).
8. DeRose, replying to Blackson, grants that he did not consider the option of SSI, explaining that he had “long assumed that SSI can’t be right” (2004: 346), since it mispredicts how high-stakes attributors treat low-stakes subjects, in a way that DeRose considers “lethal to SSI” (2004: 348). See Hawthorne 2004 (espec. pp. 160–4) and Stanley 2005 for some replies, and Schaffer 2006b (§2) for further rejoinders. In any case, DeRose’s objection to SSI has nothing to do with assertion *per se*. In the main text I will focus only on considerations arising from what one may assert.
9. Point of clarification: AQ itself allows that stakes play some further role concerning what one may assert. AQ only rules out the prospect that stakes can explain *all* the contextual dependence of what one may assert. The question under discussion plays a role, and that is all the contextualist needs. Indeed, given that the question plays a role, there is a natural psychological explanation for why stakes might *seem* to play a role. Stakes play an indirect and contingent role in the questions we consider: *the more there is at stake, the less we tend to take for granted*. (See Schaffer 2006b for a critique of SSI, and arguments that stakes have just an indirect and contingent effect on knowledge attributions, by inspiring us to query a wider range of options.)
10. See Schaffer 2004 and 2006a for these arguments. For standard forms of contextualism, see Cohen 1988, DeRose 1995, Lewis 1996, and Neta 2002, *inter alia*. I leave unresolved whether the question-relative view should count as a version of contextualism, or an alternative to it (especially since it features an invariantist semantics for “knows”). That is a merely verbal question. The substantive matter is whether the three features identified in the main text are upheld.
11. Of course there are many other objections to contextualism in the literature. But Hawthorne’s objection is the only one I am aware of that directly concerns norms of assertion, and applies to KQ.
12. Thanks to Martijn Blaauw, Jessica Brown, Herman Cappelen, Andy Egan, Ram Neta, Jason Stanley, Crispin Wright, and audiences at the Free University of Amsterdam, *Linguistics and Epistemology* in Aberdeen, the *Australasian Association of Philosophy* in Armidale, the *Pacific APA* in Pasadena, and the *Assertion Workshop* in St. Andrews.

References

- Austin, J. L. 1946. "Other Minds," *Proceedings of the Aristotelian Society* 20: 149–87.
- Blaauw, Martijn 2004. *Contrastivism: Reconciling Skeptical Doubt with Ordinary Knowledge*. Dissertation at Free University of Amsterdam.
- Blackson, Thomas 2004. "An Invalid Argument for Contextualism," *Philosophy and Phenomenological Research* 68: 344–5.
- Carlson, Lauri 1983. *Dialogue Games: An Approach to Discourse Analysis*. Reidel.
- Cohen, Stewart 1988. "How to be a Fallibilist," *Philosophical Perspectives* 2: 91–123.
- DeRose, Keith 1995. "Solving the Skeptical Problem," *Philosophical Review* 104: 1–52.
- 2002. "Assertion, Knowledge, and Context," *Philosophical Review* 111: 167–203.
- 2004. "The Problem with Subject-Sensitive Invariantism," *Philosophy and Phenomenological Research* 68: 346–50.
- Dretske, Fred 1970. "Epistemic Operators," *Journal of Philosophy* 67: 1007–23.
- Hawthorne, John 2004. *Knowledge and Lotteries*. Oxford University Press.
- Ginzburg, Jonathan 1996. "The Semantics of Interrogatives," in Shalom Lappin (ed.), *The Handbook of Contemporary Semantic Theory*: 385–422. Basil Blackwell.
- Groenendijk, Jeroen, and Martijn Stokhof 1997. "Questions," in Johan F. A. K. van Benthem and Alice G. B. ter Meulen (eds.), *Handbook of Logic and Language*: 1055–1124. MIT Press.
- Higginbotham, James 1996. "The Semantics of Questions," in Shalom Lappin (ed.), *The Handbook of Contemporary Semantic Theory*: 361–383. Basil Blackwell.
- Johnsen, Bredo 2001. "Contextualist Swords, Skeptical Plowshares," *Philosophy and Phenomenological Research* 62: 385–406.
- Lackey, Jennifer 2007. "Norms of Assertion," *Nous* 41: 594–626.
- Lewis, David 1979. "Score-keeping in a Language Game," in Rainer Bauerle, Urs Egli, and Arnim von Stechow (eds.), *Semantics from Different Points of a View*: 172–87. Springer.
- 1996. "Elusive Knowledge," *Australasian Journal of Philosophy* 74: 549–67.
- Morton, Adam and Antti Karjalainen 2003. "Contrastive Knowledge," *Philosophical Explorations* 6: 74–89.
- Neta, Ram 2002. "*S* knows that *p*," *Nous* 36: 663–81.
- Roberts, Craige 2004. "Context in Dynamic Interpretation," in Laurence R. Horn and Gregory Ward (eds.), *The Handbook of Pragmatics*: 197–220. Basil Blackwell.
- Schaffer, Jonathan 2004. "From Contextualism to Contrastivism," *Philosophical Studies* 119: 73–103.
- 2006a. "Contrastive Knowledge," *Oxford Studies in Epistemology* 1: 235–71.
- 2006b. "The Irrelevance of the Subject: Against Subject-Sensitive Invariantism," *Philosophical Studies* 127: 87–107.
- 2007. "Knowing the Answer," *Philosophy and Phenomenological Research* 75: 383–403.
- forthcoming. "Knowing the Answer Redux: Replies to Brogaard and Kallestrup," *Philosophy and Phenomenological Research*.
- Sinnott-Armstrong, Walter 2004. *Pyrrhonian Skepticism*. Oxford University Press.
- Stalnaker, Robert 1999a. "Assertion," in *Context and Content*: 78–95. Oxford University Press.
- 1999b. "On the Representation of Context," in *Context and Content*: 96–113. Oxford University Press.
- Stanley, Jason 2005. *Knowledge and Practical Interests*. Oxford University Press.
- Unger, Peter 1975. *Ignorance: A Case for Scepticism*. Oxford University Press.
- Weiner, Matthew 2005. "Must We Know What We Say?" *Philosophical Review* 114: 227–51.
- Williamson, Timothy 2000. *Knowledge and its Limits*. Oxford University Press.