Knowing the Answer to a Loaded Question

Asbjørn Steglich-Petersen
Aarhus University
filasp@hum.au.dk

Abstract
Many epistemologists have been attracted to the view that knowledge-*wh* can be reduced to knowledge-*that*. An important challenge to this, presented by Jonathan Schaffer, is the problem of 'convergent knowledge': reductive accounts imply that any two knowledge-*wh* ascriptions with identical true answers to the questions embedded in their *wh*-clauses are materially equivalent, but according to Schaffer, there are counterexamples to this equivalence. Parallel to this, Schaffer has presented a very similar argument against binary accounts of knowledge, and thereby in favor of his alternative contrastive account, relying on similar examples of apparently inequivalent knowledge ascriptions, which binary accounts treat as equivalent. In this article, I develop a unified diagnosis and solution to these problems for the reductive and binary accounts, based on a general theory of knowledge ascriptions that embed presuppositional expressions. All of Schaffer's apparent counterexamples embed presuppositional expressions, and once the effect of these is taken into account, it becomes apparent that the counterexamples depend on an illicit equivocation of contexts. Since epistemologists often rely on knowledge ascriptions embedding presuppositional expressions, the general theory of them presented here will have ramifications beyond defusing Schaffer's argument.

1. Introduction

In recent years, considerable work has been devoted to understanding the nature of knowledge-*wh* ascriptions, i.e. ascriptions of knowledge in which the complement clause is an interrogative *wh*-clause, instead of a declarative *that*-clause. Traditionally, epistemologists have devoted more attention to knowledge-*that* ascriptions, in part because it has been assumed that knowledge-*wh* could be understood in terms of, or be reduced to, knowledge-*that*. The most simple of such reductions holds that to know-*wh* is a matter of knowing that p, where p happens to be the true answer to the indirect question embedded in the *wh*-clause of the relevant knowledge-*wh* ascription. So if the game begins at 4pm, knowing *what time the game begins* is simply a matter of knowing *that the game begins at 4pm*. Reductions of this sort have the virtues of being simple and intuitively plausible, and of unifying our accounts of two kinds of knowledge ascriptions.[1]

Recently, however, Jonathan Schaffer (2007; 2009) has presented what he takes to be a decisive argument against the possibility of reducing knowledge-*wh* to knowledge-*that*. His argument relies on a series of alleged counterexamples to a

---

[1] Versions of this reductive account can be found in Hintikka (1975), Lewis (1982), Higginbotham (1996), and Stanley and Williamson (2001).

principle, which is implied by any reductive account of the above sort, namely that knowledge-*wh* ascriptions with identical true answers to their embedded questions (what Schaffer calls 'convergent' knowledge-*wh* ascriptions) are materially equivalent. Schaffer also relies on these examples to motivate an alternative, non-reductive theory of knowledge-*wh*, according to which knowledge-*wh* is question-relative in a more fundamental sense than supposed by the reductive account.

Parallel to this, Schaffer has argued against *binary* accounts of knowledge in general, i.e. the common assumption among epistemologists that knowledge is a two-place relation between a knower and a propositional content. He has instead proposed that knowledge is a three-place relation between a knower, a propositional content, and a *relevant contrast* (Schaffer 2004; 2005a; 2008. Schaffer and Knobe 2012). His main argument against the binary account, and in favor of his contrastive alternative, is very similar to his argument against reductive accounts of knowledge-*wh*: examples of intuitively inequivalent knowledge ascriptions are presented, which binary accounts treat as equivalent, while the contrastive account treat them as inequivalent.[2]

In this article, I seek to show that Schaffer's counterexamples to reductive accounts of knowledge-*wh* and binary accounts of knowledge both fail, for the same reason. Despite appearances, Schaffer does not succeed in finding examples of convergent knowledge-*wh* ascriptions that aren't materially equivalent, and he does not succeed in finding inequivalent pairs of knowledge ascriptions that binary accounts of knowledge treat as equivalent.[3]

The key observation is that all of Schaffer's examples of inequivalent knowledge ascriptions state the known content using presuppositional expressions. I argue that when knowledge ascriptions state the content of the knowledge in such a way, the knowledge ascription as a whole presupposes that the subject of the ascription knows the presuppositions carried by the content statement. This means that when evaluating the truth value of such knowledge ascriptions, we tend to assume that the person to whom the knowledge is ascribed knows the presuppositions. But if that is the case, knowledge ascriptions that *differ* in the presuppositions carried

---

[2] The contrastive account is a version of contextualism, and Schaffer thus also takes the examples to speak against invariantism, and in favor of contextualism. Since the debate over contextualism raises a number of issues that are beyond the scope of this paper, I will not pursue this here. However, if Schaffer's argument against the binary account turns out to be defective, the use of that argument against invariantism will be defective too.

[3] For an argument against Schaffer's related contrastive account of *causation*, see Steglich-Petersen (2012).

by the content statements will tend to be evaluated in *epistemically different contexts*, i.e. the contexts in which the respective presuppositions are known by the subject.[4] This affects our evaluation of the knowledge ascriptions' truth-value, but crucially, differences in evaluation caused in this way does not allow us to conclude that the relevant knowledge ascriptions are inequivalent, since that would require an illicit equivocation between the epistemically distinct contexts in which the ascriptions are evaluated. All of Schaffer's examples of seemingly inequivalent knowledge ascriptions can be demonstrated to rely on this kind of illicit context equivocation.

This makes the present defense of the reductive and binary accounts importantly different from previous defenses, both in strategy and substance. For example, Kallestrup (2009) and Brogaard (2009) both defend the reductive account of knowledge-*wh* by arguing that the reductive account is compatible with the relevant pairs of knowledge ascriptions being inequivalent. By contrast, the defense to be pursued in this article rejects that the relevant knowledge ascriptions are inequivalent. Aloni and Égré (2010) propose an analysis of the type of disjunctive knowledge-*wh* ascriptions Schaffer relies on, according to which Schaffer's examples of convergent but inequivalent knowledge-*wh* ascriptions aren't convergent in the first place. By contrast, the present account *accepts* the convergence of the relevant knowledge ascriptions, and offers the significant advantage over Aloni and Égré of generalizing to Schaffer's counterexamples to binary accounts of propositional knowledge, and to those of Schaffer's examples that don't involve disjunctive constructions, thereby providing a unified account of the problematic cases. A further significant difference between the present proposal and that of Aloni and Égré, is that while their analysis follows Schaffer in his overall contextualist understanding of knowledge, the present analysis is compatible with invariantism about knowledge. They also do not refer the relevant contextual or domain restriction effects systematically to the presupposition facts relied upon by the present account. Gerken (2012; 2013) rejects some of the cases of inequivalence discussed in this article, in the context of defending strict invariantism from Schaffer's argument for contrastivism (which is a form of contextualism). Gerken explains the intuitive appeal of the inequivalences as the mistaken result of a focal bias, causing intuitive judgments about knowledge

---

[4] In the following, whenever I speak of someone knowing something 'in a context', I always have the subject's context in mind, and not the context of the person ascribing the knowledge. My treatment of these cases does not, therefore, commit me to any kind of contextualism about knowledge.

ascriptions to be generated by processing only a limited part of the available information, namely the part that is in focus, because of limitations of cognitive resources. By contrast, the present defense does not rely on supposing any bias or limitation of processed information on behalf of those evaluating the relevant knowledge ascriptions, but instead explains the intuitive judgments as the result of reliance on information encoded by presuppositional expressions in the knowledge ascriptions.[5]

The rest of the article is structured as follows: In Section 2, I outline Schaffer's argument against reductive accounts of knowledge-*wh*; in Section 3, I outline Schaffer's argument against binary accounts of knowledge; in Section 4, I develop and motivate a theory of knowledge ascriptions involving presuppositional expressions, and make a prediction on how we tend to evaluate them based on that theory; in Section 5, I apply the theory and its prediction to show that Schaffer's alleged counterexamples rely on an illicit equivocation between contexts of evaluation; in Section 6, I show that Schaffer's examples actually *depend* on the tendency predicted by the theory for eliciting the problematic intuitions, making a denial of the theory dialectically ineffective for Schaffer; in Section 7, I conclude that Schaffer's counterexamples are ineffective.

2. The Problem of Convergent Knowledge

As indicated, the problem of convergent knowledge is a problem arising for accounts of knowledge-*wh* which seek to reduce this kind of knowledge to knowledge-*that*. Although there are differences in detail between existing reductive accounts of knowledge-*wh*, according to Schaffer all reductive accounts entail the following reductive principle, which I shall follow Shaffer in referring to as [RED] in the following (2007: 386):

[RED]   *s* knows-*wh* iff *s* knows that *p*, where *p* is the true answer to the indirect question *Q* of the *wh*-clause.

---

[5] The idea that presupposition figures into an account of knowledge-*wh* was briefly suggested by Kent Bach in a short unpublished commentary on Jonathan Schaffer's 'Knowing the Answer', presented at a 2005 conference. The suggestion is not developed further, however.

Schaffer seeks to refute reductive accounts of knowledge-*wh* by showing that there are counterexamples to [RED], rendering the principle and thus reductive accounts entailing it false. Explaining this argument requires defining some key terms.

First, call any set of questions 'convergent' if and only if they have the same true answer. Second, call any pair or set of knowledge-*wh* ascriptions 'convergent knowledge ascriptions' if and only if their interrogative *wh*-clauses denote convergent questions. For example (to use Schaffer's own cases), if we suppose that George W. Bush is the person on television, the questions 'Is it Bush or Janet Jackson on television?' and 'Is it Bush or Will Ferrell on television?' are convergent questions, since they have the same true answer, namely 'Bush is on television'. This also renders the following two knowledge-*wh* ascriptions convergent:

[1] Sally knows whether Bush or Janet Jackson is on television.

[2] Sally knows whether Bush or Will Ferrell is on television.

[RED] entails that convergent knowledge-*wh* ascriptions are materially equivalent. If any two knowledge-*wh* ascriptions to the same subject *s* embed questions with the same true answer *p*, [RED] entails that both of these knowledge-*wh* ascriptions will be true if and only if *s* knows that *p*, thereby making them materially equivalent whenever they are convergent. It is important to emphasize that on Schaffer's use, 'material equivalence' simply means identity in truth-value, and [RED] implies that the relevant knowledge ascriptions are equivalent in this sense *only when convergent*. That convergent knowledge ascriptions are 'materially equivalent' does thus *not* imply that they express identical propositions, share truth-conditions, or some other similarly demanding condition. Whenever I speak of 'equivalence' in the following, it is intended in this undemanding sense.

The problem is, according to Schaffer, that not all convergent knowledge ascriptions are materially equivalent, even in this undemanding sense. Among the counterexamples mentioned by Schaffer are [1] and [2] above. Here's Schaffer:

> [[1] and[2]], for instance, are clearly inequivalent. Knowing whether Bush or Janet Jackson is
> on television is a relatively *easy* task. The question of whether it is Bush or Janet Jackson is a
> relatively easy question. Virtually anyone (with decent vision and minimal cultural
> background) can know whether it is Bush or Janet Jackson. In contrast, knowing whether

Bush or Will Ferrell is on television is a relatively *hard* task. If the impersonation is good enough, the question of whether it is Bush or Will Ferrell may be a rather hard question. Perhaps only the first lady Laura Bush will be able to tell the difference. So one might know whether Bush or Janet Jackson is on television, but fail to know whether Bush or Will Ferrell is on television. (2007: 387)

As a further counterexample, Schaffer asks us to suppose that there is a goldfinch in the garden, rendering the following knowledge-*wh* ascriptions convergent:

[3]   Sally knows whether there is a goldfinch in the garden, or a raven.

[4]   Sally knows whether there is a goldfinch in the garden, or a canary.

Again, [RED] treats [3] and [4] as materially equivalent, since according to [RED] they are both true if and only if Sally knows that there is a goldfinch in the garden. But according to Schaffer, they are *not* equivalent:

[[3] and [4]] are clearly inequivalent. Knowing whether there is a goldfinch in the garden or a raven is a relatively *easy* task of *bird-identification*. Virtually anyone (with decent vision and minimal expertise) can know whether there is a goldfinch in the garden or a raven. In contrast, knowing whether there is a goldfinch in the garden or a canary is a *harder* task of bird identification. Perhaps only an expert birder will be able to tell the difference. (2007: 388)

If Schaffer is right that the pairs [1]/[2] and [3]/[4] are materially inequivalent even when convergent, i.e. are such that there are contexts in which they differ in truth-value even while sharing their true answer, [RED], and any reductive account of knowledge-*wh* committed to [RED], fails. Schaffer says little beyond what is contained in the above quotations to motivate or explain why these pairs are to be regarded as inequivalent. But the intuition that they are inequivalent does seem relatively compelling. It is clearly easier to tell Bush apart from Janet Jackson than from Will Ferrell, and it is likewise clearly easier to tell goldfinches apart from ravens than from canaries. So it does seem easier for [1] and [3] to obtain, than for [2] and [4] to obtain. This seems to provide a strong case for inequivalence, and therefore against reductive accounts.

3. A Similar Problem for Binary Accounts of Knowledge

Parallel to this, Schaffer has put forward a general theory of knowledge (encompassing both knowledge-*wh* and propositional knowledge), according to which knowledge is a *three-place* relation between a subject, a proposition, and a relevant contrast (*s* knows that *p* rather than *q*), instead of a *two-place* relation between a subject and a propositional content (*s* knows that *p*), as traditionally supposed (Schaffer 2004; 2005a; 2008. Schaffer and Knobe 2012).[6] His main argument in favor of the contrastive account, and against traditional binary accounts, is similar to his argument against the reductive theory of knowledge-*wh*: Sets of seemingly inequivalent knowledge ascriptions are presented, which the binary account treats as equivalent, while the alternative contrastive account correctly predicts and explains their inequivalence. This argument was first presented in Schaffer (2008); a slightly altered version of the argument is supported by experimental data in Schaffer and Knobe (2012). I shall briefly present both versions.

In his (2008), Schaffer focuses on pairs of cases such as the following, which differ only in how the questions concerning the target knowledge states are phrased, specifically in which lexical items and grammatical constructions they use for encoding *contrasts*:

Rather:

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

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

(Schaffer 2008: 236)

Here, the contrasts in the questions are encoded using *rather than* constructions. The other cases presented in (Schaffer 2008) encode the contrasts using *whether*, the *who/what* distinction, cleft constructions, and non-neutral intonation, in similar ways. Schaffer points out that we tend to answer 'Yes' to the first question, but 'No' to the second. So given the very same background information about the relevant case, i.e.

---

[6] The contrastive account was, to the best of my knowledge, first proposed by Antti Karjalainen and Adam Morton in their (2003).

that Mary has stolen a bicycle from a toy store, and that the detective has found Mary's fingerprints at the crime scene, we tend to agree that the detective knows that Mary rather than Peter stole the bicycle, but tend to disagree that the detective knows that Mary stole the bicycle rather than the wagon. The same holds for the other cases in which the questions differ only in the way the contrasts are encoded.

On the face of it, this result spells trouble for the binary account of knowledge. On that account, the knowledge states under question in (e) and (f), and the rest of the similar question-pairs considered by Schaffer, are identical in reducing to the state of knowing *that Mary stole the bicycle*. This means that the binary account is committed to regarding the knowledge ascriptions in question as sharing truth-value with the relevant binary knowledge ascription, and therefore with each other. Adopting the terminology of the previous section, the binary account will therefore regard the knowledge ascriptions in question as materially *equivalent*, while the contrastive account will regard them as materially *inequivalent*.[7] So if Schaffer is right that case (e) and (f) (and the rest of the presented pairs) represents identical contexts in which the knowledge ascriptions differ in truth-value, (e) and (f) together constitute a counterexample to the binary account, and evidence in favor of the contrastive account.

In their (2012), Schaffer and Knobe bolster this argument using more detailed cases supported by experimental data. Experimental subjects were presented with the following vignette, named the *jewel thief vignette*, similar in content to the background information given in the (2008) version of the argument:

> Last night, Peter robbed the jewelry store. He smashed the window, forced open the locked safe, and stole the rubies inside. But Peter forgot to ware gloves. He also forgot the security camera.
>
> Today, Mary the detective has been called to the scene to investigate. So far she has the following evidence. She has been told that there was a theft, she has found and identified Peter's fingerprints on the safe, and she has seen and recognized Peter on the security video, filmed in the act of forcing open the safe. She has no further information.
> (Schaffer and Knobe 2012: 689)

---

[7] In the case of knowledge-*wh*, the material equivalence was restricted to *convergent* knowledge ascriptions. Since convergence is a property of knowledge-*wh* only, this restriction does not apply to the material equivalence to which the binary account of knowledge is committed.

One group of experimental subjects were then asked whether they agreed or disagreed with [5], stating a knowledge ascription using *rather* to encode a contrast concerning the identity of the thief, while another group were asked whether they agreed or disagreed with [6], stating a knowledge ascription using *rather* to encode a contrast concerning the identity of the thing stolen:

[5]   Mary knows that Peter rather than anyone else stole the rubies.

[6]   Mary knows that Peter stole the rubies rather than anything else.

Subjects were asked to rate the statements on a scale from 1 ('disagree') to 7 ('agree'), and it was found that they tended agree with [5] (mean rating: 4.6), and disagree with [6] (mean rating: 3.1), to a statistically significant degree.

Schaffer and Knobe also tested for a contrast effect in responses to statements using the *who/what* distinction to encode the same contrasts as before. One group of experimental subjects were asked whether they agreed or disagreed with [7], stating a knowledge ascription using a *knows who* construction to encode a contrast concerning the identity of the thief, while another group were asked whether they agreed or disagreed with [8], stating a knowledge ascription using a *knows what* construction to encode a contrast concerning the identity of the thing stolen:

[7]   Mary knows who stole the rubies.

[8]   Mary knows what Peter stole.

Again, subjects were asked to rate the statements on a scale from 1 ('disagree') to 7 ('agree'). It was found that they tended agree with [7] (mean rating: 4.91), and disagree with [8] (mean rating: 2.62), to a statistically significant degree.

These results seem to speak against a binary understanding of the knowledge ascribed in [5]-[8], since on that account, all of the ascribed knowledge states reduce to knowing that *Peter stole the rubies*. Adopting the terminology of the previous sections, the binary account will regard the knowledge ascriptions in question as *equivalent*, while the contrastive account will regard them as *inequivalent*. So if Schaffer is right that [5]/[6], and [7]/[8] respectively represent pairs of knowledge ascriptions that differ in truth-value in the very same context, these pairs constitute

counterexamples to the binary account and evidence in favor of the contrastive account.

Schaffer and Knobe also considered knowledge ascriptions that do not feature any explicit encoding of contrasts, such as: 'Mary knows that Peter stole the rubies'. On the face of it, such ascriptions pose a challenge to the contrastive account, according to which *all* knowledge ascriptions are at least implicitly contrastive. But Schaffer and Knobe suggest that when a contrast is not explicitly specified in the knowledge ascription itself, people look to information from the conversational context to fill in the relevant contrast. In order to test this hypothesis, Schaffer and Knobe conducted a third study featuring a set of stimuli mirroring the two earlier experiments, except that the contrast condition was manipulated using conversational context alone, by manipulating the question under discussion. All experimental subjects received the same jewel thief vignette used in the two prior studies. Experimental subjects in the thief-contrast condition then read:

> Everyone is now asking the big question: Who stole the rubies? The news reporter is about to write a story about Mary. He is wondering if Mary now knows who stole the rubies. He writes: "Mary now knows that Peter stole the rubies."

> Please tell us whether you agree or disagree with the news reporter's claim, "Mary now knows that Peter stole the rubies."

Experimental subjects in the jewel-contrast condition instead read:

> Everyone is now asking the big question: What did Peter steal? The news reporter is about to write a story about Mary. He is wondering if Mary now knows what Peter stole. He writes: "Mary now knows that Peter stole the rubies."

> Please tell us whether you agree or disagree with the news reporter's claim, "Mary now knows that Peter stole the rubies."

Again, participants were asked to rate the statements on a scale from 1 ('disagree') to 7 ('agree'). It was found that they tended agree with the knowledge ascription in the thief-contrast condition (mean rating: 5.24), and disagree with the knowledge ascription in the jewel-contrast condition (mean rating: 2.97), to a statistically significant degree. The results thus continued the pattern from the first two

experiments. To Schaffer and Knobe, this suggests that the knowledge ascriptions in the two conditions are inequivalent despite being linguistically identical, because they are implicitly supplied with different contrasts by the conversational context.

In the following I shall seek to show that the apparent inequivalences acting as counterexamples to the reductive account of knowledge-*wh* and binary accounts of knowledge in general are illusions stemming from equivocating contexts in which the respective knowledge ascriptions are true. For all that Schaffer has shown, the relevant pairs of knowledge ascriptions are equivalent after all, despite appearances. The key to realizing this, is to understand the effect of presuppositional expressions in the content statement of knowledge ascriptions. I shall now turn to develop a theory of this.

4. Knowledge Ascriptions and Presuppositions

It is a familiar observation that many lexical items and grammatical constructions can trigger presuppositions of the utterances they are part of. Although there is substantial disagreement among linguists about the exact nature of presuppositions, and the mechanisms by which they are triggered, they can be loosely but uncontroversially characterized as implicit factual assumptions whose truth is taken for granted by the speaker of the presupposition carrying utterance, and must be antecedently accepted, or be 'common ground', among the other conversational participants in order for the utterance to be felicitous in a context.[8][9]

So-called 'presupposition-triggers' include *definite noun phrases*, which presuppose reference or unique reference of the description ('The present King of France is bald' presupposes that there is a King of France); *factive verbs*, which presuppose the truth of their object ('Sally regrets that the bill wasn't passed'

---

[8] By 'felicitous', I shall mean something like *situationally appropriate*, in a way that goes beyond being true, grammatically well-formed, etc. To give another common example of felicity conditions, we tend to find S asking H about p felicitous in a context only if S doesn't know whether p, S wants to know whether p, and S believes that H might know whether p.

[9] This conception of presuppositions was developed in detail by Robert Stalnaker (1970; 1974). For a recent survey, see Beaver (2007). There is disagreement on whether to regard the presuppositions of an utterance as psychological states of the speaker, i.e. as background beliefs whose truth the speaker takes for granted in making the presupposition carrying utterance (as the above loose characterization suggests), or whether to understand presuppositions semantically as propositions that must be true in order for the presupposition-carrying utterance to be felicitous. Nothing hinges on this debate for the present purposes, as long as the presuppositions of an utterance are such that it, under normal circumstances, can be reasonably inferred from the utterance that the speaker believes the presuppositions. Both accounts of presuppositions are compatible with this common criterion for identifying presuppositions.

presupposes that the bill wasn't passed); *aspectual verbs* ('I stopped smoking last year' presupposes that I have smoked in the past); *cleft constructions* ('It was Susan that stole the bicycle' presupposes that someone stole a bicycle); and *questions* ('Which of the houses did you buy?' presupposes that the person being asked the question bought a house).

Presupposition differs from *entailment* insofar as presuppositions are preserved under negation of the presupposition-carrying utterance: both 'Peter has stopped smoking' and 'Peter has not stopped smoking' presuppose that Peter has smoked. This doesn't mean that it is strictly incompatible with the negation of a presupposition-carrying utterance, that the presupposition is false. That Peter has not stopped smoking is strictly speaking compatible with Peter never having smoked. But it would be highly infelicitous to utter 'Peter has not stopped smoking' while being aware that Peter has never smoked.

Presupposition also differs from, and is stronger than, pragmatic *implicature* insofar as the presuppositions of an utterance cannot be cancelled: it would be highly infelicitous, not to say absurd, to state 'Peter has stopped smoking' and then immediately add 'I don't mean to suggest that Peter has ever smoked'. Two common 'tests' for presuppositions is thus for them to be preserved under negation and be uncancellable. I will return to these tests later on.[10]

Presupposition failure occurs when the proposition presupposed to be true is in fact false. There is a venerable debate on the effect this has on the presupposition carrying utterance: does it make the utterance *false*, as Russell (1905) famously claimed, or does it make the utterance

 defective in the sense of lacking a truth-value, as Strawson (1950) suggested? I shall not contribute to this debate presently, but will assume what seems to be the currently dominant view, namely that presupposition failure *sometimes* causes the utterance to be false, and *sometimes* causes it to be without truth-value, depending, among other things, on the importance of the presupposition to the topic of discourse. For example, if we are discussing who attended yesterday's exhibition, and I utter 'The exhibition was visited yesterday by the king of France', we'd be inclined to deem the sentence *falsified* by it wrongly presupposition that there is a king of France. But an utterance that clearly takes the king of France as its primary topic of discourse,

---

[10] There are other tests for presuppositions, such as projection in questions, antecedents of conditionals, *etc*., but I shall restrict myself to the two most common tests here.

such as the classic 'The King of France is bald', we'd be inclined to find defective in the stronger sense of lacking a truth-value.[11]

Of particular importance for our purposes is the *pragmatic role* commonly played by presupposition. As mentioned above, presuppositions can be regarded as implicit factual assumptions whose truth is taken for granted by the speaker of the presupposition carrying statement, and must be antecedently accepted, or be 'common ground', among the other conversational participants in order for the utterance to be felicitous in a context. This allows speakers to rely on presupposition triggers to pragmatically implicate which parts of the uttered content the speaker regards as common ground, and which the speaker regards as introducing new information to the conversational context. For example, if I utter *Peter has stopped smoking*, rather than uttering, say, *Peter used to smoke, but has stopped*, I signal to the other conversational participants that I regard it as common ground that Peter used to smoke, and that the new information to the conversational context conveyed by the utterance is the information about Peter having stopped.

But what about cases where the presuppositions of an utterance are *not* common ground? There are two main types of cases to consider. In cases where the presuppositions are not common ground because the other conversational participants antecedently *believe them to be false*, the other conversational participants must typically reject the utterance and challenge the presuppositions, rather than evaluate or answer the utterance directly. For example, if I believe that Peter has never smoked, I would not directly agree or disagree with the statement *Peter has stopped smoking*, but instead challenge the presupposition that Peter has smoked. More importantly for our purposes, in cases where the presuppositions are not common ground because the other conversational participants *didn't know of them*, conversational participants tend to *accommodate* the presuppositions, unless they have independent grounds for doubting them, and focus their attention on evaluating what the speaker has signaled to be the new information conveyed by the utterance. So if I utter *Peter has stopped smoking*, other conversational participants will tend to take for granted or accommodate that Peter used to smoke and focus attention on him having stopped, even if they didn't initially have any beliefs about Peter having smoked (unless, of course, they have independent grounds for doubting the presupposition). This phenomenon of presupposition accommodation is so natural

---

[11] For defenses of this view, see e.g. Strawson (1964), von Fintel (2004) and Yablo (2006).

that speakers often rely on presupposition triggers to make implicit assertions about matters that they know the other conversational participants to be ignorant about. For example, it would be quite natural for someone to report to a complete stranger that 'My husband is arriving shortly', without first informing the stranger of the presupposition carried by that statement, i.e. that the speaker has a husband, and simply rely on the stranger to accommodate this from the statement.[12]

It is important to emphasize that although presupposition triggers can be abused to manipulate conversational participants into an unwarranted or biased focus on a limited part of the information conveyed by an utterance, causing performance errors of various kinds in evaluating the presupposition carrying utterance, such rhetorical tricks actually trade on it normally being *warranted* to restrict one's attention as a response to cooperative speakers' use of presupposition triggers to pragmatically implicate that the presuppositions can be treated as common ground. Cooperative speakers observe the Maxim of Quality (Grice 1989), urging speakers not to say what they lack evidence for or don't believe, and there is no reason to suppose that this Maxim doesn't also extends to what is merely presupposed. So when cooperative speakers make presuppositions, it is normally warranted for the other conversational participants to accommodate those presuppositions if they don't already believe them, and focus attention on what is indicated to be the new relevant information conveyed by the utterance.

A significant clue to resolving Schaffer's challenges for the reductive account of knowledge-*wh* and binary accounts of knowledge comes from noting that *all* of Schaffer's examples encode the content of the ascribed knowledge using presupposition-triggering expressions. The seemingly inequivalent knowledge-*wh* ascriptions encode their contents using *interrogatives* and disjunctive *whether* constructions, which are well-known presupposition triggers. And the knowledge ascriptions acting as counterexamples to binary accounts of knowledge encode their contents using *who/what* distinctions, cleft constructions, non-neutral intonation, and disjunctive constructions using *whether* and *rather*. All of these are well-established devices for encoding presuppositions. This warrants investigation into the hypothesis that the use of presupposition triggers in encoding the contents of the problematic knowledge ascriptions is what accounts for our intuitions of inequivalence. I shall

---

[12] This account of presupposition accommodation is often attributed to Lewis (1979). For a detailed discussion of this complex topic, see also Beaver and Zeevat (2007).

now introduce a principle that I will rely on to pursue that hypothesis. I will first explain what the principle says, and then move on to motivate it.

Call the following principle the *Knowledge Ascription Presupposition Principle*, or [KAP] for short:

[KAP]   When a knowledge ascription states the content of the ascribed knowledge in a presupposition carrying way, the knowledge ascription as a whole carries the presupposition that the subject of the knowledge ascription knows the presupposition carried the statement of the content.

Since ascriptions of both propositional knowledge and knowledge-*wh* can state the content of the ascribed knowledge in a presupposition carrying way, [KAP] applies to both types of knowledge ascriptions.

When applied to ascriptions of propositional knowledge, [KAP] states that when a knowledge ascription *S knows that p* states the content *p* in a way that carries presupposition *q*, the knowledge ascription as a whole carries the presupposition that *S* knows that *q*. Likewise, when applied to ascriptions of knowledge-*wh*, [KAP] states that when a knowledge ascription *S knows-wh* states the content indicated by the interrogative *wh*-clause in a way that carries presupposition *q*, the knowledge ascription as a whole carries the presupposition that *S* knows that *q*.

It is important to note that [KAP] concerns *two levels* of presuppositions, and their interaction: the presuppositions carried by the statement of the *content* of the knowledge state ascribed, and the presuppositions carried by the knowledge ascription as a whole. An example will help clarify this. Consider the following knowledge ascription:

[9]   Sally knows that Peter has stopped smoking.

The content of the knowledge state ascribed by [9] is stated in a presupposition carrying way. Specifically, the statement of the content in the form *Peter has stopped smoking* carries the presupposition that Peter used to smoke. According to [KAP], this means that [9] as a whole carries the presupposition that Sally knows that Peter used to smoke. Likewise for the knowledge-*wh* version: According to [KAP], the knowledge-*wh* ascription *Sally knows whether Peter has stopped smoking* carries the presupposition that Sally knows that Peter used to smoke.

If that is what the principle says, what motivates it? To begin with, we can test [KAP] using the standard tests for presuppositions mentioned above. The first test is that of not being cancellable. [KAP] clearly predicts presuppositions that pass this test. For example, it would not be possible without absurdity to state [9], and then immediately add, for example, *I don't mean to suggest that Sally knows that Peter used to smoke*. The second test is preservation under negation. We must ask, for example, whether the posited presupposition of [9] is preserved under negation of [9] as in [10]:

[10]   Sally does not know that Peter has stopped smoking.

One way of deciding this, is to reapply the cancellability test to [10]. And again, it seems that [KAP] succeeds in predicting presuppositions that pass the test. It does seem highly infelicitous to assert [10], and then immediately add *I don't mean to suggest that Sally knows that Peter used to smoke*.

In addition to predicting presuppositions that pass the standard tests, [KAP] can be motivated on theoretical grounds as a particular instance of general and widely accepted principles for when presuppositions 'project' from sentences embedded under various operators (a phenomenon known as 'presupposition projection'). Karttunen (1974) distinguishes between sentential operators that inherit the presuppositions from the sentences they embed (known as 'holes'), and sentential operators that don't inherit presuppositions in this way (known as 'plugs'). Factive verbs are among the operators that inherit the presuppositions of their embedded sentences (in addition to presupposing the truth of the *asserted* content of the embedded sentences). For examples, 'Sally regrets that Peter stopped smoking', 'Sally is surprised that Peter stopped smoking', 'Sally remembers that Peter stopped smoking', etc., not only presuppose the truth of the embedded sentence 'Peter stopped smoking', but also inherit the presupposition of the embedded sentence, namely that Peter used to smoke. Among the factive verbs with this property is 'knows'. So sentences embedding other sentences within the operator 'knows' inherit the presuppositions of the embedded sentences. For example, 'Sally knows that Peter stopped smoking' presupposes that Peter used to smoke.

This, of course, is not by itself enough to motivate [KAP]. But we can get closer to a motivation by noting how presuppositions are projected from sentences

within the scope of *non-factive* propositional attitude verbs, such as 'believe', 'think', 'expect', 'fear', 'intend', 'suspect', 'assume', 'wants', 'hope', etc. As noted by Karttunen (1974), these are all 'plugs' in the sense defined above.[13] For example, 'Sally believes that Peter stopped smoking' doesn't carry the presupposition that Peter used to smoke. However, Karttunen noted that verbs such as these are subject to a different rule: if a non-factive propositional attitude verb embed a sentence carrying presupposition *p*, the sentence as a whole presuppose that the subject of the propositional attitude *believes* the presupposition of the embedded sentence. For example, 'Sally believes that Peter stopped smoking', 'Sally fears that Peter stopped smoking', etc., presuppose that Sally *believes* the presupposition of the embedded sentence that Peter used to smoke. If we add to this the uncontroversial assumption that ascribing knowledge that *p* involves ascribing belief that *p* (because knowledge requires belief), we can conclude that if a knowledge ascription embed a sentence carrying the presupposition *p*, the knowledge ascription as a whole carries the presupposition that the subject believes that *p*. For example, 'Sally knows that Peter has stopped smoking' presupposes that Sally believes that Peter used to smoke (this is easy to confirm using the standard tests).

Putting the above together, we can derive that knowledge ascriptions embedding presupposition-carrying sentences presuppose that the embedded presuppositions are both *true* and *believed* by the subject of the knowledge ascription. 'Sally knows that Peter has stopped smoking' presupposes that it is *true* that Peter used to smoke, and that Sally *believes* that Peter used to smoke. Admittedly, this is not quite the same as deriving [KAP], which claims that the embedded presuppositions are not only presupposed to be *true* and *believed*, but that they are *known*. But in conjunction with the independent plausibility of [KAP] as demonstrated by it passing the standard tests, I submit that this is sufficient as motivation for the principle for the present purposes.[14]

---

[13] See also Heim (1992) for further discussion.
[14] Dretske (1970) famously denied that the knowledge operator 'penetrates' to presuppositions, but I take this interesting claim to be sufficiently controversial not to require separate discussion in this context.

If [KAP] is right, and it can be assumed that conversational participants tend to accommodate presuppositions made by speakers as described above, the principle gives rise to the following prediction:[15]

[PREDICTION] When evaluating knowledge ascriptions, we tend to assume that the subject of the ascription knows any presuppositions carried by the statement of the content of the ascribed knowledge state, regardless of whether the knowledge ascription as a whole is true or false, and evaluate whether the knowledge ascription is true or false on that assumption.

When applied to ascriptions of propositional knowledge, [PREDICTION] says that when evaluating a knowledge ascription *S knows that p*, where the content *p* is stated in a way that carries presupposition *q*, we tend to assume that *S* knows that *q* regardless of whether the knowledge ascription *S knows that p* as a whole is true or false, and evaluate the truth of the knowledge ascription on that assumption. Likewise, when applied to ascriptions of knowledge-*wh*, [PREDICTION] says that when evaluating a knowledge ascription *S knows-wh*, where the content indicated by the interrogative *wh*-clause is stated in a way that carries presupposition *q*, we tend to assume that *S* knows that *q* regardless of whether the knowledge ascription *S knows-wh* as a whole is true or false, and evaluate the truth of the knowledge ascription on that assumption.

Again, an example will help to clarify this. Suppose that we are asked to evaluate [9], based on being told of some specific body of evidence available to the subject of the knowledge ascription. The content of the knowledge state ascribed in [9] is stated in a way that carries the presupposition that Sally knows that Peter has smoked. Given this presupposition, [PREDICTION] predicts that when evaluating [9], we will tend to assume that Sally knows that Peter has smoked, and evaluate whether she knows that has stopped smoking, on that assumption. This means that we will tend to focus on Sally's evidence concerning whether or not Peter presently smokes (Does he smell like cigarette smoke? Does he step outside at regular intervals? Is he prone to coughing? *etc*.), rather than evidence concerning Peter having smoked in the past (Has she seen Peter smoking in the past? Has he told her

---

[15] Arguably, this prediction could be made on the basis of principles weaker than [KAP]. For example, if it were merely pragmatically implied rather than presupposed that the subject knows the presuppositions of the content statement, the tendency described by [PREDICTION] would arguably still be in effect.

that he used to smoke? *etc*.). If we find that her evidence is adequate for knowing that Peter does not presently smoke, [PREDICTION] thus predicts that we will tend to find [9] true, independently of any assessment of her evidence concerning Peter having smoked in the past. In this example, this means that [PREDICTION] predicts that we in many cases will find it true that Sally knows that Peter has stopped smoking, even if we know of no evidence suggesting that Sally knows that Peter has smoked in the past, because we instead simply accommodate the presupposition of Sally knowing this. On the face of it, this seems a plausible prediction.

It is important to emphasize that in typical cases, the tendency to take knowledge of presuppositions for granted described by [PREDICTION] does not rest on a bias or limitation in our evaluation of knowledge ascriptions. On the contrary, in typical conversational contexts, taking such knowledge for granted will be justified by reliance on the implication made by the cooperative speaker in phrasing the content of the relevant knowledge state in a presuppositional way.[16] Another important point to emphasize is that [PREDICTION] is merely a statement about what we *tend* to do when evaluating this kind of presuppositional knowledge ascriptions. In particular [PREDICTION] does not imply that it would be impossible, or even always irrelevant, when evaluating the relevant kind of knowledge ascriptions, to also consider the evidence the subject has for the presuppositions carried by the content of the ascribed knowledge. For example, if one has independent reason to suppose that Sally *doesn't* know that Peter has smoked in the past, it would be both possible and relevant to reject [9] as relying on the false presupposition of Sally knowing that Peter has smoked in the past. Finally, it should be mentioned that [PREDICTION] is silent as to whether the tendency it describes is conscious and controlled, or unconscious and automatic. In many cases, accommodation of presuppositions is clearly unconscious and automatic, as when we accommodate from 'My husband is arriving shortly' that the speaker has a husband. This is also what makes the use of presuppositions effective as a rhetorical devise for focusing the attention of conversational participants on selected parts of the content communicated. But [PREDICTION] is also compatible with the tendency sometimes being the result of conscious reasoning about presuppositions.

---

[16] As we shall see, this marks an important difference between the present explanation of Schaffer's examples, and explanations based on hypothesized biases, cognitive limitations, or performance errors, such as that proposed by Gerken (2012; 2013).

5. Defusing the Counterexamples

To see the relevance of [KAP] to our evaluation of Schaffer's puzzling pairs of knowledge ascriptions, we can begin by substituting those ascriptions into [PREDICTION], which is based on [KAP]. In order to do so, we must first determine what presuppositions these knowledge ascriptions make. Consider first [1], i.e. *Sally knows whether it is Bush or Janet Jackson on TV*. The content of the knowledge ascribed in [1] is given by the embedded question *Is it Bush or Janet Jackson on TV?* This question carries the presupposition that either Bush or Janet Jackson, but not both, is on TV. This assumes, of course, that Schaffer in this and similar examples intends the disjunction in the *exclusive* sense, i.e. as being true if and only if exactly one of the disjuncts is true. But if the disjunction were intended in an *inclusive* sense, i.e. as being true if one or both of the disjuncts are true, Schaffer's discussion of the relative difficulty of discriminating between Bush and Janet Jackson, and Bush and Will Ferrell, respectively, would not apply to the examples.[17] Given this, we can substitute into [PREDICTION] to arrive at the following instance of that prediction:

[PREDICTION 1]   When evaluating whether *Sally knows whether it is Bush or Janet Jackson on TV*, we tend to assume that Sally knows that either Bush or Janet Jackson, but not both, is on TV, regardless of whether the knowledge ascription as a whole is true or false, and evaluate whether the knowledge ascription is true or false on that assumption.

This becomes interesting when compared to the analogous instance of [PREDICTION] for [2], i.e. *Sally knows whether it is Bush or Will Ferrell on TV*. This ascription embeds the question *Is it Bush or Will Ferrell on TV?*, triggering the presupposition that it is either Bush or Will Ferrell, but not both, on TV. This yields the following:

[PREDICTION 2]   When evaluating whether *Sally knows whether it is Bush or Will Ferrell on TV*, we tend to assume that Sally knows that either Bush or Will Ferrell, but

---

[17] See Aloni et al (2013) for a detailed discussion of the different possible interpretations this type of disjunctive knowledge-*whether* ascriptions. As they point out, there are more possible interpretations that the one relied on here, and in some contexts, some of these alternative interpretations may be more natural. As mentioned above, however, it seems clear from Schaffer's discussion concerning the relative difficulty of the relevant identification tasks that it is the exclusive sense he has in mind (see the quoted passages), and as long as this is how we *tend* to interpret the case, the prediction holds.

not both, is on TV, regardless of whether the knowledge ascription as a whole is true or false, and evaluate whether the knowledge ascription is true or false on that assumption.

The difference between these two predictions is important, since in order to show that [1] and [2] are *inequivalent* in the relevant sense (i.e. of different truth-value), even when convergent, we must find at least one single context in which Bush is on TV (thus rendering them convergent), and [1] and [2] differ in truth value. But [PREDICTION 1] and [PREDICTION 2] together suggest that when we evaluate [1] and [2], we actually tend to consider two epistemically *different* contexts: one in which knowledge that one of either Bush or Janet Jackson is on TV is taken for granted, and one in which knowledge that one of either Bush or Will Ferrell is on TV is taken for granted. So tending to judge [1] true, and [2] false, does not by itself support the conclusion that the two ascriptions are inequivalent in the relevant sense. For all we know, this tendency is explained by us judging [1] true in one context, and [2] false in another, thus leaving it an open question whether the two knowledge ascriptions share truth-value when considered in *identical* contexts.

But comparing [PREDICTION 1] and [PREDICTION 2] also explains why [1] and [2] *appear* inequivalent, in a way that is compatible with them both reducing to knowledge *that it is Bush on TV*, thus rendering them in fact equivalent whenever they are convergent. Schaffer's explanation, recall, turned on it being 'easier' to tell Bush apart from Janet Jackson than from Will Ferrell. This led Schaffer to suppose that one might know whether it is Bush or Janet Jackson on TV, but at the very same time fail to know whether it is Bush or Will Ferrell on TV, even when it in fact is Bush on TV. But given [PREDICTION 1] and [PREDICTION 2], there is a more straightforward explanation. Note first that it is much easier to come to know *that it is Bush on TV* in a context where it can be taken as prior knowledge that it is Bush or Janet Jackson on TV, than it is in a context where it can be taken as prior knowledge that it is Bush or Will Ferrell on TV. If one already knows that the person one is looking at on TV is either Bush or Janet Jackson, it takes relatively minimal further perceptual evidence to determine that it is Bush (evidence that it is a male person would suffice, for example). But if one merely knows that one is looking at either Bush or Will Ferrell, it will take relatively more evidence to get into a position where one can determine that it is Bush one is looking at. But crucially, none of this implies

that there could be *single epistemic context* in which it is Bush on TV, such that Sally knows whether it is Bush or Janet Jackson on TV, but doesn't know whether it is Bush or Will Ferrell on TV, since this difference in relative epistemic difficulty is a difference that depends on the knowledge being acquired in *two epistemically different contexts*. To conclude that [1] and [2] are inequivalent even when convergent would thus be to equivocate between the two epistemically different contexts we tend to consider when evaluating [1] and [2].

We can lend further support to this conclusion by considering two contexts that vary with respect to whether the relevant presuppositions of [1] and [2] are satisfied by the subject, but are identical with respect to the further evidence available to the subject, where this evidence is designed to let the subject discriminate between Bush and Janet Jackson, but not between Bush and Will Ferrell. Here's the first such context:

[Context 1]   Presuppositions of [1] are satisfied by Sally, but those of [2] are not: Sally has been told by a very reliable producer of the television program she is watching, that in a moment either Bush or Janet Jackson, but not both, will appear on television.

                          Further evidence available to Sally: Casting a quick glance at the television screen, Sally sees that it is a swaggering grey haired man appearing.

For our purposes, there are three relevant questions about the knowledge Sally has, or is in a position to have, in this context. First we must ask whether Sally is in a position to know that it is Bush on television. And here, the answer seems a clear 'yes'. Sally knows from the producer that the person on television is either Bush or Janet Jackson, and seeing that it is a swaggering grey haired man is sufficient to rule out Janet Jackson, in which case she may infer that the person is Bush. It seems equally clear that Sally is in a position to know whether it is Bush or Janet Jackson that is on television, by reasoning identical to the above. But the important question is whether she is also in a position to know whether it is Bush or Will Ferrell on television. On the face of it, this question should strike us as suffering from presupposition failure. After all, the context has been designed to satisfy the presuppositions of [1], but not those of [2]. But there is a line of reasoning available to Sally from the presuppositional knowledge of [1] together with the further evidence, which allows her to satisfy the presuppositions of [2] as well: If Sally

knows that the person on television is either Bush or Janet Jackson, and knows from the additional evidence that the person isn't Jackson, she can conclude that the person is Bush. But in that case, she is also in a position to know that the person is either Bush or Ferrell, thus satisfying the presuppositions of [2]. So although [2] should strike us as somewhat unnatural in the context, it does not formally suffer from presupposition failure. With this in place, it should be immediately obvious that [2] is true in the context. If Sally knows that it is either Bush or Janet Jackson, she can rule out in that it is Will Ferrell. This allows her to conclude on the basis of the additional evidence, that the person is Bush. But in that case, she will be in a position to know whether it is Bush or Will Ferrell, thus making [2] true. This result is interesting insofar as the additional evidence available in [Context 1] is of the exact sort that we should expect *not* to allow Sally to know whether it is Bush or Will Ferrell, according to Schaffer's interpretation of the intuition we have of this being relatively 'harder' to know.

   With that in mind, we can consider the next context.

[Context 2]   Presuppositions of [2] are satisfied by Sally, but those of [1] are not: Sally has been told by a very reliable producer of the television program she is watching, that in a moment either Bush or Will Ferrell, but not both, will appear on television.

       Further evidence available to Sally: Casting a quick glance at the television screen, Sally sees that it is a swaggering grey haired man appearing.

Once again, we can ask whether Sally is in a position to know what we considered above. First, does Sally know that it is Bush on television? Here the answer seems a clear 'no'. Knowing in advance only that the person is either Bush or Will Ferrell, the further evidence does not allow Sally to determine that it is Bush. Neither does she know whether it is Bush or Will Ferrell on television, for the same reason. But crucially, it doesn't seem like Sally knows whether it is Bush or Janet Jackson either, even if the further evidence available to Sally enables her to discriminate Bush and Jackson. There are two ways of coming to this conclusion. First we might simply observe that if the evidence doesn't allow Sally to know that it is Bush and doesn't allow her to know that it is Janet Jackson, we should not be inclined to say that it allows her to know whether it is Bush or Janet Jackson. The evidence will, of course, put her in a position to know that it *isn't* Janet Jackson. But this still leaves open the

possibility that it is Will Ferrell, and doesn't, therefore, allow her to know whether it is Bush or Janet Jackson.[18] We can arrive at the same conclusion by noticing that the presuppositions of [1] are unsatisfied, and that there isn't a line of reasoning available to Sally that will satisfy them, on the basis of the presuppositions of [2] and the further evidence. If Sally merely knows in advance that it is either Bush or Will Ferrell, being able to exclude that it is Janet Jackson will not enable her to know that it is either Bush or Janet Jackson. So [Context 2] doesn't satisfy the presuppositions of [1] in the way [Context 1] satisfied those of [2]. So [1] is false in [Context 2]. This is interesting insofar as it is in contexts with evidential backgrounds such as this that we should expect Schaffer's example of 'easy' knowledge to obtain.

The upshot of this is that when we take prior knowledge of presuppositions into account, [1] and [2] retain material equivalence when convergent, even across contexts in which the evidence is of the sort that, according to Schaffer, should lead us to expect that they will differ in truth-value. It would be impossible, of course, to consider *all* of the possible contexts in which [1] and [2] are convergent to check for material equivalence. But the two contexts considered above do seem telling of a general pattern, since the additional evidence introduced in them is designed to elicit the kind of inequivalence judgment that Schaffer trades on.

The above procedure can easily be extended to account for the intuitions of inequivalence in the rest of Schaffer's puzzling pairs of knowledge ascriptions. Since examples [3] and [4] are knowledge-*wh* ascriptions structurally identical to [1] and [2], I will leave it to the reader to apply the procedure to those, and instead proceed to show how a similar strategy can be used to defuse the counterexamples to binary accounts of knowledge. Since [7] and [8] also concerns knowledge-*wh* ascriptions, making it obvious how to apply the above procedure, I shall focus on [5] and [6] and the cases depending on clues from the conversational context.

Statement [5] was that *Mary knows that Peter rather than anyone else stole the rubies*. The *rather than* construction used for encoding the content of the ascribed knowledge is a presupposition trigger, causing the statement of the content to carry the presupposition that it was rubies that were stolen.[19] Given this, we can substitute into [PREDICTION] to arrive at the following instance of that prediction:

---

[18] Aloni and Égré (2010) argue this as well.

[19] As above, there is room for alternative interpretations of the presuppositions. But as long as there is a tendency towards interpreting them as I have done here, the prediction will hold. It is consistent with

24

[PREDICTION 5]   When evaluating whether *Mary knows that Peter rather than anyone else stole the rubies*, we tend to assume that Mary knows that its was rubies that were stolen, regardless of whether the knowledge ascription as a whole is true or false, and evaluate whether the knowledge ascription is true or false on that assumption.

This becomes interesting when compared to the analogous instance of [PREDICTION] for [6], i.e. *Mary knows that Peter stole the rubies rather than anything else*. The *rather than* construction causes the statement of the content to carry the presupposition that it was Peter that stole something. This yields the following:

[PREDICTION 6]   When evaluating whether *Mary knows that Peter stole the rubies rather than anything else*, we tend to assume that Mary knows that it was Peter that stole something, regardless of whether the knowledge ascription as a whole is true or false, and evaluate whether the knowledge ascription is true or false on that assumption.

These predictions together suggest that when we evaluate [5] and [6], we tend to evaluate them in two epistemically *different* contexts: one in which Mary knowing that it was rubies that were stolen is taken for granted, and one in which Mary knowing that it was Peter that stole something is taken for granted. But if the two knowledge ascriptions are evaluated in epistemically different contexts, a difference in truth-value doesn't constitute a counterexample to their equivalence.

The difference between the two predictions also explains the difference in the degree to which experimental subjects agreed with [5] and [6], in a way that is compatible with them both reducing to the non-contrastive, binary knowledge state of Mary knowing that *Peter stole the rubies*. To realize that, we can again consider two contexts which differ with respect to prior knowledge of the presuppositions, but are identical with respect to the further evidence available to Mary the detective. If [5] and [6], along with the non-contrastive binary state of Mary knowing that Peter stole the rubies, share truth-value across these contexts, this is strong evidence that the

the experimental results and the explanation adopted here that some experimental subjects interpret the presuppositions differently.

intuition of inequivalence is generated by context equivocation. In this case, the further evidence available to Mary is provided by the jewel thief vignette:

[Context 5]   Presuppositions of [5] are known by Mary, but those of [6] are not: Mary has been informed by a reliable source prior to arriving to the crime scene that it was rubies that were stolen, but has not been told that Peter stole them.

Further evidence available to Mary: Mary has found and identified Peter's fingerprints on the safe, and she has seen and recognized Peter on the security video, filmed in the act of forcing open the safe.

We can now assess the truth-value of our three target knowledge ascriptions. Starting with [5], *Mary knows that Peter rather than anyone else stole the rubies*, it should be clear that this is true. But interestingly, it should be equally clear that [6] is true, i.e. *Mary knows that Peter stole the rubies rather than anything else*. After all, in this context Mary already knew that it was rubies that were stolen, so if the further evidence enables her to identify the thief as Peter, she knows that Peter stole the rubies rather than anything else. But in that case, the non-contrastive binary knowledge ascription, *Mary knows that Peter stole the rubies* is true as well, of course. So the three target ascriptions share truth-value so far. What about the reverse context?

[Context 6]   Presuppositions of [6] are known by Mary, but those of [5] are not: Mary has been informed by a reliable source prior to arriving to the crime scene that it was Peter that stole something, but has not been told that it was rubies that he stole.

Further evidence available to Mary: Mary has found and identified Peter's fingerprints on the safe, and she has seen and recognized Peter on the security video, filmed in the act of forcing open the safe.

It should be clear that in this context, all three knowledge ascriptions are *false*. It should be immediately obvious that Mary doesn't know that Peter stole the rubies rather than anything else, and equally obvious that Mary doesn't know that Peter stole the rubies. But crucially, it also seems clear that Mary doesn't know that Peter rather than anyone else stole the rubies, once it is made explicit that Mary doesn't know in advance that it was rubies that were stolen. It would seem, thus, that equivalence is

upheld across the two contexts, suggesting that the intuition of inequivalence stems from equivocating between contexts in which knowledge of the respective presuppositions is assumed.

This leaves it to be explained why similar intuitions of inequivalence can be elicited in cases where only the conversational context is manipulated, and the knowledge ascription itself remains constant, as indicated by Schaffer and Knobe's third study. In this study, recall, participants were asked to evaluate whether *Mary knows that Peter stole the rubies* on the basis of the jewel thief vignette, and it was found that answers differed depending on the additions to the conversational context that participants were supplied with. Participants that were supplied with additions to the conversational context making salient the question *Who stole the rubies?* tended to agree that *Mary knows that Peter stole the rubies*, whereas participants that were supplied with additions making salient the question *What did Peter steal?* tended to disagree with the very same knowledge ascription. Since [PREDICTION] concerns how we tend to evaluate a certain kind of presuppositional knowledge ascriptions, it cannot be applied directly to explain this case. But very similar considerations about question-presupposition can be invoked explain it. Note first that the question *Who stole the rubies?* carries the presupposition that it was rubies that were stolen. So if "everyone is asking the big question: Who stole the rubies?", as the vignette states, it is presupposed, i.e. supposed to be common ground by everyone, that it was rubies that were stolen. As the vignette is set up, it is very natural to include Mary in the context in which the relevant question is being asked, in which case it is also very natural to assume that Mary is part of the common ground presupposed by the question. But if Mary is assumed to share this presuppositional knowledge, it is no wonder that she is judged to know that Peter stole the rubies – after all, she knows that it was rubies that were stolen, and she has evidence identifying Peter as the thief. If, on the other hand, she is merely assumed to share the knowledge presupposed by the question *What did Peter steal?*, i.e. that it was Peter that stole something, it becomes implausible on the basis of the vignette to judge that Mary knows that Peter stole the rubies. So the difference in verdict between the two cases can be explained by the different presuppositional knowledge Mary is assumed to share by the questions that are salient in the two contexts.

A possible objection to this explanation of the inequivalence-intuition in the case of [5] and [6], and the study manipulating the conversational context, is that the

jewel thief vignette says explicitly that Mary has no further information than that stated in the vignette. It therefore violates this limitation to suppose that we actually tend to evaluate [5] and [6] on the assumption of *further* knowledge being available to Mary, i.e. the knowledge accommodated from the statement's respective presuppositions. But it is clearly not possible for Schaffer and Knobe to control what information we rely upon in evaluating statements such as [5] and [6], just by telling us that we are not to supply any further information to the scenario than what is given in the vignette. In fact, the vignette already depends on respondents supplying a large amount of background assumptions to the scenario, including assumptions about the information and background knowledge available to Mary, such as knowledge about the nature of fingerprint evidence and video recordings. So unless it is explicitly specified whether some particular piece of information can or cannot be relied upon, it is doubtful that it can be controlled whether respondents will supply such information independently or as a result of presupposition accommodation. Even if it were made more explicit in the vignette that Mary doesn't have any evidence that it was rubies that were stolen, it is plausible that the tendency described by [PREDICTION] would still be in effect, making a difference in verdict between e.g. [5] and [6] expectable as a result of that. As argued in Section 4, presupposition triggers have the effect of focusing attention of selected parts of the content conveyed by the statement they are part of, often in an unconscious and automatic way. Normally, such focus is warranted as an accommodating response to cooperative speakers' use of presupposition triggers to indicate what can be treated as common ground. But even when the focus is unwarranted, because it removes attention from information that is relevant to our evaluation of the statement, or when there is a conflict between what is presupposed by the statement and the background information that the statement is to be evaluated in response to, the tendency described by [PREDICTION] may plausibly in many cases be the same. So the objection that the above explanation of the seeming inequivalences violates the stated limitation on the information available to Mary looks less than compelling.

6. Disregarding Knowledge of Presuppositions

I have argued that by assuming [KAP], it is possible to predict and explain the appearance of inequivalence between the problematic pairs of knowledge ascriptions,

in a way that preserves their actual equivalence, thus saving the reductive/binary accounts. But it can also be shown that if we consciously *refrain* from assuming knowledge of presuppositions, despite the tendency described by [PREDICTION], it is no longer possible to elicit the intuitions of inequivalence. The intuitions that Schaffer's counterexamples trade on thus *depend* on subjects evaluating the examples in the manner described by [PREDICTION]. This means that denying [KAP] and its predictions would be dialectically ineffective in a defense of Schaffer's counterexamples.

To see how the appearance of inequivalence depends on the tendency described by [PREDICTION], it will be sufficient to focus on a single case, such as that provided by [3] and [4]:

[3]   Sally knows whether there is a goldfinch in the garden, or a raven.

[4]   Sally knows whether there is a goldfinch in the garden, or a canary.

Assume that [3] and [4] are convergent, i.e. that the bird in the garden is a goldfinch. The question we must then ask is if there is any conceivable evidence that Sally could acquire, which would enable her to know whether there is a goldfinch in the garden, or a raven (thus making [3] true), but *not* enable her to know whether there is a goldfinch in the garden, or a canary (thus making [4] true), *if we consciously refrain* from assuming that Sally satisfies the presuppositions of the two knowledge ascriptions.

Schaffer says little on what such evidence would look like, but claims that what makes it 'easier' to know whether it is a goldfinch in the garden, or a raven, compared to knowing whether it is a goldfinch in the garden, or a canary, is that it is easier to discriminate between goldfinches and ravens, than between goldfinches and canaries. So let us consider evidence that *conclusively* discriminates between goldfinches and ravens, for example by conclusively ruling out that the bird in the garden is a raven, but doesn't discriminate between goldfinches and canaries. We can imagine, for example, that Sally is told that the bird in the garden isn't black (I assume here that all ravens are black). This evidence would certainly not suffice for [4], but would it suffice for [3]? I think that once we explicitly disregard prior knowledge of presuppositions, and remind ourselves that *all* the evidence Sally has is

that of being told that the bird isn't black, is should be clear that Sally *doesn't* know whether it is a goldfinch in the garden, or a raven. After all, this evidence is compatible with the bird being of a great deal of other species as well – any species of non-black birds. So even evidence that conclusively rules out ravens doesn't suffice for [3]. It doesn't, therefore, provide us with a kind of evidence that makes [3] true and [4] false.

Perhaps what is needed is evidence that is even more discriminatory, i.e. evidence that doesn't merely rule out ravens, but even more alternatives to the bird being a goldfinch. Suppose, for example, that Susan is given conclusive evidence that the bird in the garden belongs to the genus Carduelis, which, in addition to goldfinches, includes just five other species of birds, such as greenfinches and redpolls. Would this alone suffice for knowing whether it is a goldfinch in the garden, or a raven, once again reminding ourselves to disregard knowledge of presuppositions? It seems not. At most, it would suffice for knowing whether the bird is a member of the Carduelis genus, or a raven. In fact, this evidence alone doesn't even suffice for knowledge taking quite remote alternatives, such as knowledge whether it is a goldfinch in the garden, or an elephant. Again, it would at most be sufficient for knowing whether it is a member of the Carduelis genus, or an elephant. And needless to say, it doesn't suffice for [4], either. So even *very* discriminatory evidence doesn't provide a case of [3] and [4] coming apart in truth-value. By extending the same reasoning, it seems that evidence ruling out *even more* alternatives, short of leaving goldfinches the only viable alternative left, would also be insufficient for [3], and thus insufficient for [4], in the absence of knowledge of presuppositions.

The reasonable conclusion to draw from these considerations is not, of course, the unduly skeptical one that evidence in favor of the bird being a goldfinch must be infallible in order to suffice for [3]. Rather, the conclusion we should draw is that the only evidence that would suffice for knowing whether it is a goldfinch in the garden, or a raven, is evidence in favor of it being a goldfinch, that we would normally regard as sufficient for knowing *that it is a goldfinch*, without mention of any alternatives. But in that case, the evidence needed for knowing whether it is a goldfinch in the garden, or a raven, would *also* suffice for knowing whether it is a goldfinch or a canary, in which case there would be no difference between the evidence needed for the two to obtain at all.

The upshot of this is that the appearance of inequivalence disappears when we consciously refrain from assuming knowledge of presuppositions in the way predicted by [KAP]. This means that any appearance of inequivalence there might be when considering these pairs must stem from our tendency to assume knowledge of presuppositions. But if that is the case, the appearance of inequivalence is illusory: it depends on equivocating between contexts in which knowledge of the presuppositions of the respective knowledge ascriptions are assumed to hold.

7. Conclusion

I have argued that the problem of convergent knowledge for reductive accounts of knowledge-*wh*, and the problem for binary accounts of knowledge, both fail for the same reason. They both depend on examples of apparently inequivalent knowledge ascriptions, which the reductive and binary accounts treat as equivalent. But for all we know, these knowledge ascriptions are equivalent after all.

This also means that Schaffer's main argument in favor his alternative question-relative account of knowledge-*wh* becomes defective. His main argument in favor of this account is its ability to predict and explain why pairs such as [1] and [2], and [3] and [4] are inequivalent even when convergent. However, since these pairs turn out to be equivalent when convergent, predicting and explaining their inequivalence becomes a liability rather than an argument in favor of the question-relative account. Likewise for Schaffer's main argument in favor of the contrastive account of knowledge, which also relies on the inequivalence of knowledge ascriptions, that on inspection turned out to be equivalent.

Many arguments in epistemology depend on eliciting intuitive responses to knowledge ascriptions that are stated in subtly different ways, often using highly sensitive linguistic devises, such as presupposition triggers. If [KAP] and the prediction derived from it are correct, Schaffer's arguments are examples of how such devises can mislead us when comparing differently stated knowledge ascriptions. There are surely many other interesting effects of presupposition triggers in the context of knowledge ascriptions awaiting further exploration.[20]

References

Aloni, Maria and Paul Égré. 2010. "Alternative Questions and Knowledge Attributions," *The Philosophical Quarterly* 60: 1-27.

Aloni, Maria, Paul Égré, and Tikitu de Jager (2013). "Knowing Whether A or B," *Synthese* 190: 2595-2621.

Beaver, David I. 1997. "Presupposition," in Johan van Bentham and Alice Ter Muelen (ed.), *Handbook of Logic and Language*: 939-1008. Elsevier.

Beaver, David I. and Henk Zeevat 2007. "Accommodation," in Gillian Ramchand and Charles Reiss (eds.), *Oxford Handbook of Linguistic Interfaces*. Oxford University Press.

Brogaard, Berit 2009. "What Mary Did Yesterday: Reflections on Knowledge-Wh," *Philosophy and Phenomenological Research* 78: 439-67.

Gerken, Mikkel 2012. "On the Cognitive Bases of Knowledge Ascriptions," in Mikkel Gerken and Jessica Brown (eds.), *Knowledge Ascriptions*. Oxford University Press.

Gerken, Mikkel (2013). "Epistemic Focal Bias," *Australasian Journal of Philosophy* 91: 41-61.

Grice, H. Paul (1989). "Logic and Conversation," in *Studies in the Way of Words*: 1-144. Harvard University Press.

Heim, Irene 1992. "Presupposition Projection and the Semantics of Attitude Verbs," *Journal of Semantics* 9: 183-221.

Higginbotham, James 1996. "The Semantics of Questions," in Shalom Lappin (ed.), *The Handbook of Comtemporary Semantic Theory*: 361-83. Oxford University Press.

Hintikka, Jaakko 1975. "Different Constructions in Terms of the Basic Epistemological Verbs: A Survey of Some Problems and Proposals," *The Intensions of Intentionality and Other New Models for Modalities*: 1-25. Dordrecht.

Kallestrup, Jesper 2009. "Knowledge-Wh and the Problem of Convergent Knowledge," *Philosophy and Phenomenological Research* 78: 468-76.

Karjalainen, Antti and Adam Morton 2003. "Contrastive Knowledge," *Philosophical Explorations* 6: 74-89.

Karttunen, Lauri 1974. "Presupposition and Linguistic Context," *Theoretical Linguistics* 1: 181-94.

Lewis, David 1979. "Scorekeeping in a Language Game," *Journal of Philosophical Logic* 8: 339-59.

Lewis, David 1982. " Whether Report," in Tom Pauli (ed.), *Philosophical Essays Dedicated to Lennart Åqvist on his Fiftieth Birthday*: 194-206. University of Uppsala.

Russell, Bertrand 1905. "On Denoting," *Mind* 14: 479-93.

Schaffer, Jonathan 2004. "From Contextualism to Contrastivism," *Philosophical Studies* 119: 73-104.

Schaffer, Jonathan 2005a. "Contrastive Knowledge," *Oxford Studies in Epistemology* 1: 235-71.

Schaffer, Jonathan 2005b. "Contrastive Causation," *Philosophical Review* 114: 327-58.

Schaffer, Jonathan 2007. "Knowing the Answer," *Philosophy and Phenomenological Research* 75: 383-403.

Schaffer, Jonathan 2008. "The Contrast-sensitivity of Knowledge Ascriptions," *Social Epistemology* 22: 235-45.

Schaffer, Jonathan 2009. "Knowing the Answer Redux: Replies to Brogaard and Kallestrup," *Philosophy and Phenomenological Research* 78: 477-500.

Schaffer, Jonathan and Joshua Knobe 2012. "Contrastive Knowledge Surveyed," *Noûs* 46: 675-708.

Stalnaker, Robert C. 1970. "Pragmatics," *Synthese* 22: 272-89.

Stalnaker, Robert C. 1974. "Pragmatic Presupposition," in Munitz et al (eds.), *Semantics and Philosophy: Essays*. New York University Press.

Stanley, Jason and Timothy Williamson 2001. "Knowing How," *Journal of Philosophy* 98: 411-44.

Steglich-Petersen, Asbjørn 2012. "Against the Contrastive Account of Singular Causation," *British Journal for the Philosophy of Science* 63: 115-43.

Strawson, Peter 1950. "On Referring," *Mind* 59: 320-44.

Strawson, Peter 1964. "Identifying Reference and Truth-Values," *Theoria* 30: 96-118.

von Fintel, Kai 2004. "Would you believe it? The king of France is back!

Presuppositions and truth-value intuitions," in Reimer, M. and Bezuidenhout, A. (eds.), *Descriptions and Beyond*, Oxford University Press.

Yablo, Stephen 2006. "Non-Catastrophic Presupposition Failure," Thomas, J.J. and Byrne, A. (eds.), *Content and Modality: Themes from the Philosophy of Robert Stalnaker*, Oxford University Press.