

University Press Scholarship Online

Oxford Scholarship Online



Knowledge and Presuppositions

Michael Blome-Tillmann

Print publication date: 2014

Print ISBN-13: 9780199686087

Published to Oxford Scholarship Online: August 2014

DOI: 10.1093/acprof:oso/9780199686087.001.0001

Knowledge and Presuppositions

Michael Blome-Tillmann

DOI:10.1093/acprof:oso/9780199686087.003.0002

Abstract and Keywords

This chapter introduces the author's account of the semantics of 'knowledge' attributions: *Presuppositional Epistemic Contextualism* (PEC), according to which the content of 'knows' in a context *C* is determined, in part, by what is pragmatically presupposed in *C*. After developing the Stalnakerian notion of a *pragmatic presupposition*, the chapter makes progress towards a solution of one of the major problems for relevant alternatives versions of contextualism: it provides an account of what it means for a possibility (or an alternative) to be epistemically relevant at a context while at the same time avoiding the objection that 'knowledge' becomes—as Lewis puts it—overly 'elusive'. As a result, this chapter offers a response to the familiar objection to contextualism that, without a precise account of relevance, epistemic contextualism remains unacceptably occult.

Keywords: knowledge, context, contextualism, epistemic contextualism, David Lewis, Robert Stalnaker, pragmatic presupposition, context-sensitivity, presuppositional epistemic contextualism

1.1 Epistemic Contextualism

Let me begin with a brief characterization of epistemic contextualism (EC). Informally speaking, EC is the view that there are both contexts in which it is difficult to satisfy the predicate 'know' and contexts in which it is considerably easier to do so. In other words, EC is—roughly and intuitively speaking—the view that there are contexts with low standards for the satisfaction of 'know' and contexts with high standards for the satisfaction of 'know', the relevant 'epistemic standards' being determined by certain contextual factors, such as the practical goals, interests and—as I shall argue later—the presuppositions of the conversational participants. To illustrate this view, epistemic contextualists typically compare 'know' with gradable adjectives such as 'tall' or 'flat': just as who counts as 'tall' in one context of ascription does not necessarily do so in another, who counts as 'knowing *p*' in one context of ascription does not necessarily do so in another. EC is thus a semantic, or as is often said, a linguistic view, namely the view that 'know' is a context-sensitive expression. 'Know' is a predicate that has—as Kaplan (1989) puts it—an unstable character insofar as its content may vary with the conversational context of ascription.¹ (p.12)

As will be familiar to most readers, the main evidence for EC derives from our intuitions about the truth-values of certain 'knowledge'-ascriptions. Consider, by way of example, Keith DeRose's (1992) *Bank Case*, presumably the most widely discussed example in support of EC in the literature:

Low Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Realizing that it isn't very important that their paychecks are deposited right away, Hannah says,

'I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning.'

High Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, 'I guess you're right. I don't know that the bank will be open tomorrow.'²

Our intuitions concerning *Low Stakes* are that Hannah speaks truly when she self-ascribes 'knowledge'. In *High Stakes*, however, our intuitions are reversed: in *High Stakes* our intuitions are that Hannah speaks truly when denying that she 'knows that the bank will be open on Saturday'. EC aims to account for these intuitions by claiming that it is more difficult to satisfy 'knows' in the context of *High Stakes* than it is in the context of *Low Stakes*: even though Hannah is in exactly the same epistemic position towards the proposition that the bank will be open on Saturday in both cases, she nevertheless satisfies 'knows' in *Low Stakes* but not so in *High Stakes*. More evidence or a stronger epistemic position is required for Hannah to satisfy 'knows' in *High Stakes* than in *Low Stakes*.³

Examples such as the *Bank Case* have attracted a large amount of critical attention in recent years. In particular, a variety of non-contextualist (p.13) theories have been proposed that are meant to account for the data from the *Bank Case* without appealing to semantic context-sensitivity. *Subject-Sensitive Invariantism*, for instance, claims that knowledge itself—rather than 'knowledge'-attributions—is sensitive to certain factors of the subject's rather than the ascriber's context, and a view that I have elsewhere called *Pragmatic Invariantism* claims that the data from the *Bank Case* is best accounted for by postulating conversational implicatures or related pragmatic phenomena.⁴ Yet further

recently popular views to be mentioned here are *Epistemic Relativism*, according to which ‘knowledge’-attributions are assessment-sensitive, and, of course, the formerly default invariantist view—*Moderate Insensitive Invariantism*—according to which Hannah is simply wrong when denying that she ‘knows’ in *High Stakes*.⁵ Finally, it should also be noted at this point that while each of the views just mentioned takes the data emerging from the above cases to be theoretically relevant and in demand of an explanation, those very data have, more recently, been met with increasing scepticism by ‘experimental philosophers’.⁶

In this monograph, however, I shall largely ignore the views just mentioned and the methodological criticisms of examples such as the Bank Case by ‘experimental philosophers’. While extremely interesting, these views will have to be dealt with on a different occasion in the detail they undoubtedly deserve. Let me, therefore, emphasize again that the primary goal of this monograph is not to deliver a final verdict on EC—a goal that would require a comparative evaluation of all the competing views in the area. Rather, the goal of this monograph is to develop and expand upon a novel account of EC—an account that has not received much attention in the literature as of yet but that is nevertheless exceedingly promising and attractive.⁷

Besides their aim to account for the data from examples such as the *Bank Case*, contextualists have usually also claimed that their theories have (p.14) the explanatory resources to do some interesting philosophical work—namely, to resolve sceptical puzzles.⁸ To add this extra bit of explanatory force to their purely semantic claims about ‘know’, contextualists refer—following Stewart Cohen’s (1988) early (and ground-breaking) work on the topic—to an error-theory, according to which sceptical paradoxes arise because speakers are blind towards the context-sensitivity of epistemic terms when considering sceptical arguments.⁹ To be precise, the contextualist argues that when we are puzzled by sceptical arguments, we fail to realize that the propositions expressed by their conclusions are perfectly compatible with the propositions expressed by our everyday ‘knowledge’-claims.

This is so because, contextualists have argued, the ‘epistemic standards’ operative when sceptical scenarios are at issue are exceedingly higher than the ‘epistemic standards’ in everyday contexts. Because of this error-theory, contextualists have claimed to be able to account for both the plausibility of sceptical arguments and our intuition that our everyday ‘knowledge’-ascriptions express truths.

We shall return to the topic of contextualism and scepticism—and the viability of contextualism’s error-theory—in later chapters. For the moment, note that if the predicate ‘knows’ is in fact context-sensitive in a distinctly epistemic way—that is, in a way that allows us to account for the data from the *Bank Case* and to resolve sceptical puzzles—then the question arises as to exactly which contextual features determine the (Kaplan) content of ‘knows that *p*’ at a given context. According to traditional accounts of EC, the predicate ‘know’ is sensitive to either the salience of error or counterpossibilities at a context or to what is at stake at the context. On the view developed here, however, neither of these views has it quite right. On my view, ‘know’ is sensitive to what is *pragmatically presupposed* at a context. Of course, what we pragmatically presuppose may be influenced by what is salient: typically, if a possibility of error becomes salient in a conversation, this will change the speakers’ presuppositions. If, for instance, the possibility that Frank did not post the birthday card on time becomes salient in a normal conversation, then it will typically no longer be pragmatically presupposed in that conversation that Frank posted the card on time. Similar effects can be perceived with respect to practical stakes: if the stakes are particularly high with respect to a (p.15) proposition *p* in a conversational context *C*, then it is unlikely that *p* will be pragmatically presupposed in *C*.

The goal of this chapter is thus to be more explicit where other accounts have spoken loosely of a context’s ‘epistemic standards’—standards that are sometimes allowed to shift in a seemingly ad hoc manner to suit the theoretical goals of the contextualist.¹⁰ In other words, this chapter is concerned with developing a detailed and thorough account of the contextual mechanisms at play in the semantics of ‘knows’—an account

that is systematic, intuitively plausible, explanatory, and independently motivated. Given that—as I shall do in this book—we think of ‘knowledge’ along the lines of a contextualized relevant alternatives theory, according to which one satisfies ‘knows p ’ at context C only if one can eliminate all *alternatives* or *counterpossibilities* to p that are *epistemically relevant* in C , an ideally explicit account of the metaphorical notion of an ‘epistemic standard’ at a context and thus of the semantics of ‘knows’ more generally will provide us with an account of that very notion—that is, the notion of a relevant alternative. However, such an enterprise will presumably at some point come to its natural limits. As Hawthorne (2004a, p. 61) puts it in a slightly different context:

To put flesh on the bones of the theory [. . .], we must say more about the mechanisms by virtue of which [alternatives] get to count as relevant or irrelevant. Ideally, one would wish for some kind of epistemic recipe book that specified exactly how features of context would suffice to make a certain [alternative] relevant to a particular knowledge ascription. Nothing like that recipe book is currently in our possession, nor are we close to possessing one. Perhaps the concept of knowledge is sufficiently primitive that some analytic ambitions will inevitably be foiled, including any attempt to analyze the pertinent notion of relevance. Perhaps, in disputed areas, there will be a plethora of borderline cases. [. . .] More generally, perhaps the mechanisms by virtue of which context-dependent predicates get their extension [are] neither readily accessible to a priori reflection nor fully amenable to empirical investigation. [. . .] But whether or not such a recipe book is (p.16) ultimately obtainable, one would still hope for some kind of picture of how context contributes to extension. In the case of ‘tall’, for example, while we may have no very complete account available of how context serves to set a boundary between extension and anti-extension, we can imagine well enough some reasonably satisfying account of how context determines a comparison class and in turn an extension for a tokening of that predicate.

One would hope for something similar in the case of 'know'.

(Hawthorne 2004a, pp. 61-2)

As Hawthorne emphasizes, it is not clear how far we can go with the project of providing a recipe book of what determines the content of 'know'. After all, maybe the (Kaplan) character of 'know' is not susceptible to analysis. This is, I take it, to be expected. Nevertheless, Hawthorne is certainly spot-on in demanding some story that is revealing of the mechanisms underlying the contextual shifts underlying the Bank Case data and the contextualist's resolution of sceptical paradoxes: without such a story, contextualism is devoid of real explanatory force and must therefore fail to offer a genuine solution to our problems.

Finally, it should be noted that an account of what determines the 'epistemic standards' at a context is not only of interest to epistemic contextualists, but also—as the above quote from Hawthorne suggests—to subject-sensitive invariantists and epistemic relativists. Since these theorists agree with the contextualist that there is some sensitivity to 'epistemic standards', they will also need an account of this merely metaphorical notion. Thus, subject-sensitive invariantists as well as epistemic relativists will presumably be tempted to make use for their own purposes of the account developed here.

1.2 Knowledge and Relevant Alternatives

To begin our discussion of the semantics of 'knows', let us take a closer look at David Lewis's views on scepticism and contextualism. According to Lewis:

(L) A subject S satisfies 'knows p ' in context $C \leftrightarrow$
 S 's evidence e eliminates every $\neg p$ -world, except
for those that are properly ignored in C .¹¹

(p.17)

A few clarificatory remarks are in order. Firstly, note that on Lewis's approach our evidence e consists of the totality of our

perceptual experiences and memory states. For the purposes of this book it is instructive to consider this a stipulative definition of a technical notion that Lewis employs in developing his account of the semantics of 'knows', rather than an explication or analysis of our natural language or scientific notion of evidence.¹² I shall therefore use the term 'evidence' exclusively in this stipulative sense in what follows, until Chapter 5 where I lift this stipulation to explore in more detail the relation between knowledge and our natural language (or scientific) notion of evidence, within the framework of a contextualist account of the semantics of 'knows'. This leaves us with the issue of what it means for Lewis that a possibility w is eliminated by one's evidence. Lewis offers a straightforward account of elimination: a possibility w is eliminated by an experience or memory state if the existence of the experience or the memory state (rather than its content) conflicts with w .¹³

In addition to this definition of the satisfaction of 'know' and his account of the elimination of a possibility by one's evidence, Lewis stipulates a set of 'rules of relevance' specifying which possibilities can be properly ignored in a given context. It is this set of rules that is meant to determine how the content of 'know' is influenced by particular contextual factors. The rule doing the main explanatory work with regard to both sceptical puzzles on the one hand, and problem cases such as the *Bank Case* on the other, is what Lewis calls the *Rule of Attention*:

(RA) If w is attended to by the speakers in C , then w is not properly ignored in C .

Inspired by Stewart Cohen's (1988, pp. 106–11) earlier work on the contextual salience of error possibilities, Lewis points out that (RA) eventually boils down to the apparent triviality that 'a possibility not ignored at all is *ipso facto* not properly ignored.'¹⁴ (p.18)

Note that (L) in conjunction with (RA) allows for an elegant account of the data from the *Bank Case*. There are uneliminated possible worlds in which the bank will be closed

on Saturday—worlds that Hannah can properly ignore in *Low Stakes* but not in *High Stakes*, the relevant worlds being precisely those worlds in which the bank will be closed on Saturday because it has changed its hours recently. The reason why Hannah can properly ignore these counterpossibilities in *Low Stakes* but not in *High Stakes* has, on Lewis's account, to do with the fact that they have been attended to in *High Stakes* but not so in *Low Stakes*. Thus, according to Lewis, Hannah cannot truthfully assert 'I know *p*' in *High Stakes* because the possibility that the bank has changed its hours recently is no longer properly ignored.¹⁵

Moreover, (RA) and (L) also seem to provide the resources for a straightforward resolution of sceptical puzzles. To see this, note firstly that when confronted with sceptical arguments one inevitably attends to sceptical possibilities, for sceptical hypotheses—that is, sentences expressing sceptical possibilities—form an integral part of sceptical arguments. Thus, it follows from (RA) that any context in which one considers sceptical arguments is a context in which one does not properly ignore sceptical possibilities. Secondly, conceding that sceptical possibilities resist elimination by one's evidence, it follows from (RA) and (L) that, for all propositions *p* about the external world, one does not satisfy 'knows *p*' in contexts in which one considers sceptical arguments.¹⁶ Such contexts are, as I shall henceforth put it, *sceptical contexts*. Thirdly, note that even though Lewis's account entails that we do not satisfy 'knows *p*' in contexts in which sceptical arguments are at issue, it also entails that we often do so in quotidian contexts. In quotidian contexts we do not attend to sceptical possibilities. We can therefore properly ignore sceptical possibilities, and thus usually satisfy 'knows *p*' for various propositions *p* about the external world.¹⁷ (p.19)

Lewis's views can thus be seen as accounting for both our *Anti-Sceptical Intuitions* (ASI) and our *Sceptical Intuitions* (SI), which we may represent as follows:

(ASI) People often speak truly when they assert 'I know *p*.'

(SI) People sometimes speak truly when they assert ‘Nobody knows *p*’ in contexts in which sceptical arguments are discussed.

However, if the semantic value of ‘know’ can change in a way allowing for both (ASI) and (SI) to be true, why then are we puzzled by sceptical arguments? Lewis replies that the puzzle arises because we are often unaware of the relevant contextual shifts in the content of ‘know’. We simply do not always realize that our everyday ‘knowledge’-ascriptions express propositions that are perfectly compatible with the propositions expressed by ‘knowledge’-negations in sceptical contexts.¹⁸

However, things are not quite as straightforward as they might seem. There is a fairly obvious and widely noticed problem for Lewis’s (RA)—namely, that it makes it far too difficult to satisfy ‘know’. As Michael Williams points out:

[T] he Rule of Attention makes retaining knowledge too hard. Conceding for the present that far-fetched sceptical possibilities—brains-in-vats, demon-deceivers—resist elimination by evidence, the Rule ensures that a person’s knowledge vanishes every time such a possibility enters his head.

(Williams 2001, p. 15)

As it stands, (RA) allows the mere attendance to sceptical hypotheses in a context *C* to make it impossible to properly ignore such counterpossibilities in *C*.¹⁹ As Williams points out, however, this is too strong a view. Imagine you saw your teenage son sneaking away through the window of his room late at night. When you confront him the next morning he replies somewhat desperately, ‘How do you know I left the house? I mean, for all you know you might have dreamt it. It was late at night, wasn’t it?’ On Lewis’s account you find yourself in a context in which you have to admit to your son that you do not ‘know’ that he sneaked away at night, and this surely is not just a pity, it is rather also mistaken. Of course you ‘know’ that your son sneaked away through the window of his room last night—you saw him doing so, after all. (p.20)

Lewis's (RA) is thus too strong. However, an alternative that puts you in a more authoritative position regarding your son is easily obtained. Note that by means of (RA) Lewis exploits the contrast between ignoring a proposition and attending to it. Lewis: 'if in *this* context we are not in fact ignoring it but attending to it, then for us now it is a relevant alternative.'²⁰ However, it seems obvious that, *pace* Lewis, merely attending to—or directing one's mind towards—some possibility *w* in *C* is not enough for making it impossible to properly ignore *w* in *C* in the *epistemologically relevant sense*. The notion of ignoring I have in mind is thus not that of ignoring *w* as opposed to attending to *w*, but rather that of ignoring *w* as opposed to taking *w* seriously. On this second reading, you surely can attend to the possibility that you merely dreamt that your son sneaked out of his window last night while nevertheless ignoring this possibility in a straightforwardly practical sense; you can entertain the thought that you merely dreamt, or direct your mind towards that possibility, without taking this very possibility seriously or giving it any credence.²¹

The idea of replacing Lewis's (RA) with a rule employing the notion of taking a possibility seriously instead of merely attending to it comes to mind. If a possibility is taken seriously in a context *C*—that is, if it is among the 'live options' in *C*—then it cannot be properly ignored in *C*. However, what exactly does it mean for a possibility to be a 'live option' in a context *C*? One way to explicate the notion at issue is by means of the notion of a pragmatic presupposition: a possibility *w* is taken seriously in *C* just in case *w* is compatible with the speakers' pragmatic presuppositions in *C*. On this view, we can implement the idea that 'live options' cannot be properly ignored by means of the following *Rule of Presupposition*:

(RP) If *w* is compatible with the speakers' pragmatic presuppositions in *C*, then *w* cannot be properly ignored in *C*.²²

(p.21)

Why would we want to link the content of 'know' in *C* to the speakers' presuppositions in *C* rather than to other contextual features? Firstly, note that there is a fairly intuitive sense in which it is presupposed in *Low Stakes* but not in *High Stakes* that the bank has not changed its hours recently. Thus, by adopting (RP) we have a very plausible explanation of the data from the *Bank Cases*. Moreover, note that another crucial advantage of such a move is fairly obvious: since speakers can, to a certain extent, voluntarily decide what they take seriously and which propositions they presuppose, they have, to a certain extent, voluntary control over the content of 'know' in their contexts. For instance, as long as you make clear to your son that the possibility that you dreamt seeing him sneaking out of his window is not a 'live option' in your conversation, you remain in a context in which you satisfy 'know', even though your son has drawn attention to the possibility that you might have dreamt the relevant episode.²³ Similarly, as long as the speakers in a context *C* pragmatically presuppose the negations of sceptical hypotheses, the 'epistemic standards' relevant for the evaluation of 'know' in *C* remain the standards of quotidian contexts, even though attention may have been drawn to sceptical possibilities: sceptical possibilities can still be properly ignored.²⁴ Thus, replacing Lewis's Rule (p.22) of Attention by my Rule of Presupposition avoids the above-mentioned problems pointed to by Williams.²⁵

1.3 Pragmatic Presuppositions

Under what conditions does a speaker presuppose a given proposition *p*? Of course, we have a pre-theoretical understanding of what it means to presuppose something: one presupposes *p* when one takes *p* for granted or when one assumes *p*, possibly only for the purposes of the conversation one is participating in. However, even though we have an intuitive grasp of what a presupposition is, our pre-theoretical concept is, presumably, too vague to play centre stage in a contextualist approach to the semantics of 'know'. In this section I will therefore look for an explication or sharpening of our intuitive concept that can then be shown to figure in an explanation of mostly familiar data about 'knowledge'-ascriptions and sceptical puzzles.

When discussing the notion of a presupposition in a philosophical or linguistic context, Robert Stalnaker's work on the topic comes immediately to mind. A first suggestion might thus be to adopt his rather well-developed notion of a *pragmatic presupposition* for the present purposes. And in fact—as will become obvious later—Stalnaker's notion is ideally suited for putting flesh on the skeleton of a presupposition-based EC as outlined above. Thus, if I am right, the very notion that has application in Stalnaker's accounts of linguistic phenomena as diverse as assertion, sentence presupposition, indicative conditionals, and others also plays a crucial role in the semantics of 'knowledge'-ascriptions.

What, then, is a Stalnakerian pragmatic presupposition? Before answering this question it is imperative to note that Stalnaker thinks of the notion at issue as primitive. Pragmatic presuppositions are, according to Stalnaker, propositional attitudes *sui generis* and as such unsusceptible to analysis or definition. However, even though Stalnaker intends the notion to remain ultimately undefined, he offers, throughout his work, several explications of the notion (p.23) that are meant to approximate the concept and give the reader a closer grasp of it. Stalnaker justifies this approach as follows:

It may be charged that [the concept of a pragmatic presupposition is] too unclear to be the basic [concept] of theory, but I think that this objection mistakes the role of basic concepts. It is not assumed that these notions are clear. In fact, one of the points of the theory is to clarify them. So long as certain concepts all have some intuitive content, then we can help to explicate them all by relating them to each other. The success of the theory should depend not on whether the concepts can be defined, but on whether or not it provides the machinery to define linguistic acts that seem interesting and to make conceptual distinctions that seem important. With philosophical as well as scientific theories, one may explain one's theoretical concepts, not by defining them, but by using them to account for the phenomena.

(Stalnaker 1970, p. 46; cp. also 1974, p. 50)

Bearing in mind this caveat, let us consider Stalnaker's most recent explication of the notion.

In his latest work on the topic, Stalnaker proposes a two-stage explication of the notion of a pragmatic presupposition: first, he defines what he calls *common ground* in terms of the notions of *belief* and *acceptance*, and then, in the second step, he explicates the notion of a pragmatic presupposition in terms of the notions of belief and common ground.²⁶ Here is Stalnaker's definition of the concept of common ground:

(CG) It is common ground that p in a group $G \leftrightarrow$
all members of G *accept* (for the purpose of the
conversation) that p , and all *believe* that all accept
that p , and all *believe* that all *believe* that all
accept that p , etc.²⁷

Having thus defined the notion of common ground, Stalnaker gives the following explication of a pragmatic presupposition:

(PP) x pragmatically presupposes $p \leftrightarrow x$ believes p
to be common ground.²⁸

(p.24)

Thus, according to Stalnaker's explication, a speaker pragmatically presupposes p if she believes that all members participating in her discourse accept p , believe that all accept p , believe that all believe that all accept p , etc. Pragmatic presuppositions are, accordingly, a special type of belief and, as such, a special type of propositional attitude.

I have claimed above that one advantage of (RP) over (RA) is that the participants in a conversation can decide to presuppose a proposition and thus have, to some extent, voluntary control over what 'know' expresses in their context. Stalnaker's notion of a pragmatic presupposition as just explicated, however, does not allow for voluntary presupposing: since belief is spontaneous and thus not under one's direct voluntary control, one can hardly choose to believe that a proposition p is common ground. On the basis of

(PP), presupposing is outside the realm of the voluntary. Is this a problem for my account?

Note that there are a few problems with (PP) due to the fact that it does not allow for voluntary presupposition. In fact, Stalnaker himself, in a footnote, considers the following case relating to the issue:

Foreign Language

There may in some cases be a divergence between [pragmatic] presupposition and belief [. . .]. A speaker may presume that something is common ground, even when he is only hoping that it will become common ground. Suppose I am in a country whose language I do not speak. I have no reason to think that the person I approach on the street speaks English, but I am desperate, so I try: 'Is there a public toilet nearby?' If I am lucky, it will become common [ground] that we both speak English.

(Stalnaker 2002, pp. 717, fn. 726)

On Stalnaker's view, the speaker in *Foreign Language* pragmatically presupposes the proposition that the addressee speaks—or at least understands—some English, even though he does not believe that proposition to be part of the common ground. Such an interpretation of *Foreign Language*, however, is incompatible with (PP), according to which it is a necessary condition on the speaker's presupposing that the addressee understands at least some English that he believes that proposition to be part of the common ground. Thus, if Stalnaker wants to treat *Foreign* (p.25) *Language* as a case of speaker presupposition, then (PP) needs to be amended to cover the case.

Besides Stalnaker's own case, there are other, presumably less controversial, examples causing trouble for (PP). While *Foreign Language* is—according to Stalnaker—a case in which the speaker presupposes p even though he fails to believe that p is common ground, there are also more extreme cases in which the speaker presupposes p even though he *knows* that p

is not and *will not become* common ground after the utterance.
Consider the following dialogue:

Faculty Meeting

A:

I can't come to the meeting—I have to pick up my sister from the airport.

B:

Hang on; I know that you don't have a sister. You're just making up a reason to get around the meeting!

C:

That's not true. I have a sister.

B:

No, you don't.

A:

Yes, I do! I just never told you.

C:

Relax! (to A) Independently of whether you have a sister or not, will you come to the meeting?

A:

I'm sorry, but I really won't be able to come. As I said before, I have to pick up my sister from the airport.

According to (PP), A in *Faculty Meeting* does *not* pragmatically presuppose that she has a sister when making her last assertion, for she does not believe that proposition to be common ground. After all, A knows from the course of the conversation that B does not accept and will not accept the proposition that A has a sister. However, many theorists—Stalnaker among them—take the view that sincere utterances of sentences such as 'I have to pick up my sister from the airport'—that is, sincere utterances of sentences that have semantic presuppositions—are *paradigm cases* of pragmatic speaker presupposition. On the standard view of presupposition accommodation, any speaker who asserts a sentence that semantically presupposes *p ipso facto* pragmatically presupposes *p*.²⁹ Thus, as long as we want our

account of pragmatic presuppositions to be compatible with the (p.26) standard accounts of presupposition accommodation, we need to amend (PP) for it to cover cases such as *Faculty Meeting*.³⁰

Fortunately, however, the situation is not as troublesome as it might seem, for the amendment required to cover the above cases is a relatively slight one. In fact, the key to the problem can be found in Stalnaker's earlier writings on pragmatic presuppositions. Here is a quote from (Stalnaker 1974):

I shall say that one actually does make the presuppositions that one seems to make even when one is only pretending to have the beliefs that one normally has when one makes presuppositions. Presupposing is thus not a mental attitude like believing, but is rather a linguistic disposition—a *disposition to behave in one's use of language as if one had certain beliefs*, or were making certain assumptions.³¹

From this passage we can extract the following definition of pragmatic presuppositions:

(PP*) x pragmatically presupposes p in $C \leftrightarrow x$ is disposed to behave, in her use of language, as if she believed p to be common ground in C .³²

Even though (PP*) appears promising at first sight, it might be objected that the condition it specifies is too weak. Consider the case of truthful Frank, who is always disposed to assert sentences such as (1)—that is, sentences semantically presupposing that he has a sister—simply in virtue of having a sister:

(1) I have to pick up my sister from the airport.

Since asserting sentences such as (1) seem to be behaving, in one's use of language, as if one believed it to be common ground that one has a sister, it seems (p.27) to follow that Frank constantly pragmatically presupposes that he has a sister. Even worse, generalizing from Frank's case, it seems

that speakers constantly pragmatically presuppose all sorts of propositions that fail to be part of the common ground. Does this intuitively implausible result endanger (PP*)?

To see why it does not, note that the notion of a pragmatic presupposition is a technical notion that does not necessarily coincide with our intuitions about the use of the English word 'presupposition'. Moreover, distinguishing closely between a pragmatic presupposition, which is a behavioural *disposition*, and the behavioural manifestation of a pragmatic presupposition, the implausibility of (PP*) can be explained away: truthful Frank in fact constantly pragmatically presupposes that he has a sister, but he surely does not constantly *manifest* that pragmatic presupposition.

Before moving on, let me briefly outline further the importance and fruitfulness of the notion of a pragmatic presupposition for semantic theory by means of an example. Consider, for illustration, the following pairs of sentences, whose a-sentences intuitively presuppose the propositions expressed by their corresponding b-sentences:

(2)

A:

John stopped smoking.

B:

John has been smoking.

A:

Frank knows that the book was stolen.

B:

The book was stolen.

A:

The queen of Tuvalu is at home.

B:

Tuvalu has a (unique) queen.

The propositions expressed by the b-sentences in (2) are presupposed by their corresponding a-sentences in the sense

that speakers uttering the a-sentences pragmatically presuppose rather than assert the corresponding b-propositions. In other words, Stalnaker has it that the use of the above a-sentences requires that a given pragmatic presupposition be in place in the context of utterance. When I utter 'The queen of Tuvalu is at home', my utterance requires me to pragmatically presuppose that Tuvalu has a queen: it requires me to behave (and thus to be disposed to behave), in my use of language, as if I believed it to be common ground that Tuvalu has a queen. Note that this is in fact fairly uncontroversial: pragmatically presupposing p is, after all, nothing but a way of behaving, in one's use of language, as if one believed p to be common ground. Thus, whenever one sincerely and literally utters sentences such as those in (2), one pragmatically presupposes the relevant b-propositions. Stalnaker thus has the resources to (p.28) explain a semantic property of the sentences in (2) by claiming that particular lexical items ('stop', 'know', 'the') place certain pragmatic requirements on the speaker—namely, the requirement to pragmatically presuppose the relevant b-propositions. Stalnaker's notion allows us to give an explanation of a semantic phenomenon by means of a pragmatic notion.

Besides playing a crucial role in Stalnaker's account of sentence presupposition, the notion of a pragmatic presupposition is, as I have mentioned above, central in other fields of linguistics and the philosophy of language, too. Stalnaker himself, for instance, employs it in accounting for the semantics of counterfactuals and indicative conditionals, as well as in his account of assertion. In summary, it is safe to say that the notion of a pragmatic presupposition is a familiar and important tool in the kit of the contemporary linguist and philosopher of language.

Leaving aside the details of the semantics of sentence presuppositions, note that with (PP*) we have finally arrived at an account that positions pragmatic presuppositions within the realm of the voluntary. Since one has direct voluntary control over one's behavioural dispositions, one can, on the basis of (PP*), consciously decide to presuppose a proposition p .^{33, 34} Furthermore, note that the notion of a pragmatic

presupposition thus defined stands in a very tight relationship to the notion of taking seriously. Those possibilities that are taken seriously or that are treated as the 'live options' in a (p.29) conversation are precisely those possibilities that are consistent with what is pragmatically presupposed in the corresponding context. In other words, the possibility that p is taken seriously in a conversation if the participants to that conversation are not disposed to behave, in their use of language, as if they believed p to be common ground. The notions of what is taken seriously in a conversation and of what is pragmatically presupposed in a conversation are interdefinable.³⁵ Let us now return to the topic of 'knowledge'-ascriptions.³⁶

1.4 Presuppositional Epistemic Contextualism (PEC)

Let me briefly recapitulate the discussion thus far. We have seen that the core of Lewis's approach to contextualism consists in the idea that the satisfaction of 'knows' is closely tied to the elimination of relevant counterpossibilities by one's evidence. Here is again (L):

Lewis's Semantics of 'knows p' (L)

x satisfies 'knows p ' in context $C \leftrightarrow x$'s evidence eliminates every $\neg p$ -world, except for those that are properly ignored in C .

(p.30)

My view diverges from Lewis's, however, with regard to the definition of the notion of proper ignoring. While Lewis aims to account for the context-sensitivity of 'know' by means of his *Rule of Attention*:

Rule of Attention (RA)

If w is attended to by the speakers in C , then w is not properly ignored in C .

I replaced (RA) with what I have called the *Rule of Presupposition*:

Rule of Presupposition (RP)

If w is compatible with the speakers' pragmatic presuppositions in C , then w cannot be properly ignored in C .

Furthermore, I have given substance to (RP) by explicating the notion of a pragmatic presupposition along Stalnakerian lines:

Pragmatic Presuppositions (PP*)*

x pragmatically presupposes p in $C \leftrightarrow x$ is disposed to behave, in her use of language, as if she believed p to be common ground in C .

Besides differing from Lewis's account with respect to the rules determining the set of possibilities that cannot be properly ignored at a context, the view to be defended in this book diverges from Lewis's in another important respect.³⁷ To see what I have in mind, note that Lewis's (L) is subject to rather straightforward counterexamples. In particular, note that (L) includes neither a belief condition nor a condition to the effect that one's belief must be properly based for it to qualify as 'knowledge'. Thus, a subject can satisfy the conditions specified by (L) relative to a context C while holding her belief that p on an epistemically entirely inappropriate basis—such as tealeaves reading or the testimony of an exceedingly unreliable guru. Assuming that a subject basing her belief that p in such ways does not satisfy 'knows p ' relative to any context of utterance, we need to (p.31) amend (L) by supplementing it with an additional constraint ensuring the proper basing of the subject's belief.^{38, 39}

I shall therefore, in what follows, assume the following approach to the semantics of 'knows p ':

Semantics for 'knows p ' (L)*

x satisfies 'knows p ' in context $C \leftrightarrow$

1. x 's belief that p is properly based and

2. x 's evidence eliminates all $\neg p$ -worlds, except for those that are properly ignored in C .

(L*) will be developed further in the remainder of the book, but the amendments I will propose are of a largely cosmetic nature only. The main idea underlying PEC is accurately captured by (L*), and I shall therefore, in what follows, call the conjunction of (L*), (RP), (PP*), and the remaining Lewisian rules of proper ignoring, that is the *Rule of Actuality, Resemblance, Belief, Reliability, Method, and Conservatism, Presuppositional Epistemic Contextualism* or, for short, *PEC*.⁴⁰

Before moving on to the topic of sceptical puzzles, however, it is worthwhile noting that there are further reasons to prefer an account such as (L*)—that is, an account that comprises a proper basing constraint—over Lewis's more simple (L). First, note that, according to (L), we always know what our evidence is, for our evidence eliminates, by Lewis's definition of the notions of evidence and elimination, all possibilities in which we have different evidence than we actually have.⁴¹ This is, of course, an implausible consequence that is avoided by adding the proper basing constraint in (L*): once a proper basing constraint is in place, our beliefs about what our evidence is must be properly based in order to count as 'knowledge'.

(p.32)

Secondly, note that Lewis's (L), but not my (L*), runs into serious difficulties accounting for our 'knowledge' of necessary truths. To see this, note that necessary truths are true in all possible worlds: if p is a necessary truth, then there are no $\neg p$ -worlds whatsoever. Consequently, no $\neg p$ -worlds must be eliminated by one's evidence for one to 'know p '. Thus, no matter what evidence one has, one's evidence always eliminates all $\neg p$ -worlds, and one therefore—on Lewis's approach—always satisfies 'knows p ' in any context. Naturally, such satisfaction of the predicate 'knows p ' by default (for all necessary propositions p) will strike many as unintuitive: surely there are many necessary truths that we do not know. As is familiar, however, Lewis is quite attracted to the thought:

What I choose to call ‘propositions’ are individuated coarsely, by necessary equivalence. For instance, there is only one necessary proposition. It holds in every possibility; hence in every possibility left uneliminated by *x*’s evidence, no matter who *x* may be and no matter what his evidence may be. So the necessary proposition is known always and everywhere.

(Lewis 1996, pp. 551–2; symbolism adjusted)

Of course, Lewis has a story to tell about why his view appears implausible, and—even though problematic—that story is not entirely hopeless and implausible itself.⁴² Moreover, note that this feature of Lewis’s account is generally not perceived to be a knockdown objection to his version of contextualism: epistemologists have not rejected Lewis’s views on the semantics of ‘know’ simply because they presuppose coarsely individuated propositions and therefore logical omniscience across contexts.

However this may be, a different and less controversial way to resolve the situation is by adopting the independently motivated (L*) rather than (L). To see how adding a proper basing constraint helps with the problem concerning necessary truths, consider a paradigmatic case of a belief in a necessary proposition that does not amount to ‘knowledge’. Consider irrational Lou, who believes Fermat’s Last Theorem (FLT) on the basis of tealeaves reading. Does Lou ‘know’ that FLT is true? Since there are no worlds in which FLT is false, the only condition Lou needs to satisfy in order to ‘know’ FLT is the condition that his belief is properly based. However, Lou’s belief is clearly not properly based: by assumption, Lou’s (p.33) belief is based on tealeaves reading, which does not qualify as an epistemically proper method of belief formation.⁴³ Thus, since Lou does not satisfy the first constraint in (L*), he does not satisfy ‘knows FLT’ in any context whatsoever—despite the fact his evidence eliminates every possible world in which FLT is false.

Adopting (L*) accordingly resolves the problem of logical omniscience. Given (L*), a belief in a necessary truth counts as ‘knowledge’ in a context *C* just in case the belief is properly

based. Even though this consequence of (L*) may initially seem implausible, I do not think that it is mistaken: regarding necessary truths, there are no epistemically deficient ways to believe apart from those involving elements of improper causal sustenance. Having one's belief properly based is, accordingly, sufficient for 'knowing' necessary truths.⁴⁴

Before moving on it is worthwhile noting further that (L*) does not contain as an explicit condition that p is true, if x satisfies 'knows p ' in a context C . As those familiar with Lewis's account will have noticed, this omission is intentional and does by no means commit us to the view that 'knows' is not factive. Rather, Lewis has it that (L) (and thus our condition (L*)) in conjunction with the Rule of Actuality ensures the factivity of 'knows': adding the truth condition would, therefore, be redundant. I will discuss the role of Lewis's Rule of Actuality and its relation to the factivity of 'knows' in detail in Chapter 3.7. For the moment it will suffice to note that the account proposed here is by no means intended to entail the view that 'knows' is not factive.

Finally, note that while (L*) is a biconditional, I do not mean to suggest that the principle may be understood as giving a conceptual analysis or (p.34) a reductive definition of the satisfaction of 'knows'. I take it to be rather unlikely that the enterprise of providing a definition or analysis of 'knowledge' could be successful.⁴⁵ As will become clear in later chapters, I shall therefore take the view that the notion of proper ignoring employed in (L*) cannot be reductively defined—that is, defined in terms entirely independent of 'knowledge'. Moreover, it is also worthwhile noting that even though (L*) is a version of a contextualized relevant alternatives account of knowledge, there is a view much simpler than (L*) that accurately captures the main idea underlying *Presuppositional Epistemic Contextualism* as defended throughout this book. Consider what I shall call the *Simple View*:

The Simple View (SV)

x satisfies 'knows p ' in context $C \rightarrow x$'s evidence eliminates all $\neg p$ -worlds that are compatible with what is pragmatically presupposed in C .

Note that (SV) makes a claim about the role of what is pragmatically presupposed in the context of ascription with regard to the truth-conditions of 'knowledge'-attributions: it entails a presuppositional epistemic contextualist semantics of 'knows'. However, note also that (SV) is not a biconditional: it merely explicates a necessary condition for 'knowledge'—or, to be precise, for the satisfaction of 'knows p ' in a context C . The *Simple View* accordingly does not make any claims about what 'knowledge' is, about what constitutes 'knowledge', or about how the concept of 'knowledge' is to be analyzed or explicated. Thus, if the reader does not feel attracted to the relevant alternatives account of 'knowledge' implicit in (L*), she may feel encouraged to read the claims about PEC made throughout this book as claims about the *Simple View*. As we shall see towards the end of the book, most of the philosophical and explanatory work to be done in this book can be done by the *Simple View*.

In fact, it is worth noting at this point that PEC, if understood along the lines of the *Simple View*, is perfectly compatible with a large variety of views about the nature of 'knowledge'. The *Simple View* can, for instance, be combined with a JTB account of 'knowledge' in a straightforward way:

JTB Version of PEC

x satisfies 'knows p ' in context $C \leftrightarrow$

1. x believes p ,
2. x 's belief that p is justified,
(p.35)
3. p , and
4. x 's evidence eliminates all $\neg p$ -worlds that are compatible with what is pragmatically presupposed in C .

Similarly, there are, of course, reliabilist or safety-theoretic versions of PEC:

Reliabilist Version of PEC

x satisfies 'knows p ' in context $C \leftrightarrow$

1. x believes p ,
2. x 's belief that p has been formed in a reliable way,
3. p , and
4. x 's evidence eliminates all $\neg p$ -worlds that are compatible with what is pragmatically presupposed in C .

Safety Version of PEC

x satisfies 'knows p ' in context $C \leftrightarrow$

1. x 's belief that p is safe and
2. x 's evidence eliminates all $\neg p$ -worlds that are compatible with what is pragmatically presupposed in C .⁴⁶

Thus, whatever one's favourite analysis or account of 'knowledge' is, there is a PEC-version of it that results from adding the condition explicated by (SV) to one's independently preferred account.

1.5 Concluding Remarks

Having formulated the bare bones of the view to be defended in this book let us, in the following chapter, take a closer look at how it allows for a more successful and attractive resolution of sceptical puzzles than more customary versions of contextualism.

Notes:

- (1) Some theorists (Schaffer and Szabó forthcoming) object to this characterization on the basis of the claim that 'knows' is an expression that is associated with a contextually filled argument place, and is as such not an indexical in Kaplan's sense: it is the content of the aphonic argument place rather than of the expression 'knows' that varies its content with

context. However, for reasons relating to semantic compositionality, it is not, as I argue elsewhere, obvious that an expression that is semantically associated with an aphonic contextually filled argument place does not vary its Kaplan content with context (see Blome-Tillmann ms-a).

Independently of these issues, I shall refer in this monograph to 'know' as context-sensitive and, sometimes, as having an unstable Kaplan character.

(2) For the original example see (DeRose 1992). The version quoted here is borrowed from (Stanley 2005, pp. 3–4).

(3) We shall later (Section 1.2) clarify the (at this point deliberately) vague talk of 'epistemic standards' and 'evidence' in this passage.

(4) For a discussion and defence of *Subject-Sensitive Invariantism* see (Fantl and McGrath 2002, 2009; Hawthorne 2004a; Stanley 2005). For *Pragmatic Invariantism* see (Brown 2006; Rysiew 2001, 2007) and, for critical discussion of the view, (Blome-Tillmann forthcoming).

(5) The main advocate of *Epistemic Relativism* is John MacFarlane (MacFarlane 2005, 2011); *Moderate Insensitive Invariantism* is defended explicitly and in detail by Timothy Williamson (2005a, 2005b) and, more recently, Jennifer Nagel (2007, 2008, 2010).

(6) For an interesting discussion see DeRose (2011).

(7) Of course, developing such a coherent and explanatorily powerful version of EC will in itself amount to giving an argument in favour of the view.

(8) The only exception to this claim that I am aware of is (Ludlow 2005).

(9) Cf. (Cohen 1988, p. 106; DeRose 1995, p. 40).

(10) I have in mind DeRose (2009, p. 240), who defends the view that speakers can 'select epistemic standards', and who employs this view in replying to objections to his version of EC. It should be noted that the invocation of epistemic

standards is only viable if supplemented with a more detailed and informative conception of those standards and their selection. As long as we are not told more about what epistemic standards are, how they are contextually determined, and—most crucially—how they are selected, DeRose does not really have a response to the phenomena he addresses by means of the mentioned claim. For criticism of DeRose's early conception of epistemic standards and his Rule of Sensitivity, see (Blome-Tillmann 2009a).

(11) See (Lewis 1996, p. 553). Note that Lewis speaks of 'possibilities' rather than 'worlds' in his account.

(12) Nothing in Lewis's paper suggests that he takes his definition of evidence as offering an explication or analysis of our natural language notion of evidence.

(13) (Lewis 1996, p. 553): 'I say that the uneliminated possibilities are those in which the subject's entire perceptual experience and memory are just as they actually are. There is one possibility that actually obtains (for the subject and at the time in question); call it *actuality*. Then a possibility *w* is *uneliminated* iff the subject's perceptual experience and memory in *w* exactly match his perceptual experience and memory in actuality.'

(14) (Lewis 1996, p. 559).

(15) Thus, on Lewis's view, it is not the stakes in the two cases that determine whether or not Hannah satisfies 'knows' but rather what is attended to at the relevant context, or what possibilities and counterpossibilities are salient.

(16) Note that sceptical possibilities resist elimination by one's evidence only if the contents of experiences and memories are individuated internalistically. In this chapter I shall grant the sceptic such an internalist conception of evidence.

(17) I assume that none of the other Lewisian rules of relevance such as that of Actuality, Resemblance, or Belief marks out sceptical worlds as relevant in quotidian contexts.

(18) This strategy, relying on what I have elsewhere called the phenomenon of ‘semantic blindness’, has been criticized widely, but see (Blome-Tillmann 2008) for a comprehensive defence.

(19) This is particularly absurd in cases in which a participant to a conversation attends to a sceptical possibility in their own thought only, i.e. without mentioning the possibility to other speakers.

(20) (Lewis 1996, p. 559; Lewis’s emphasis).

(21) Lewis himself canvasses a normative variant of his position, which, he acknowledges, conflicts with (RA)—namely by modifying (L) so that it ends ‘except for those possibilities which we could properly have ignored [if we hadn’t attended to them]’ (Lewis 1996, pp. 560–1). However, this normative approach effectively eliminates (RA) and thus the contextualist element from Lewis’s approach. Lewis’s normative approach is to be paired with criteria distinguishing those possibilities that one can properly ignore (or could have properly ignored) in a context from those that one cannot properly ignore (or could not have properly ignored) in a context. In what follows, I offer such criteria.

(22) Jonathan Schaffer (2004a, 2005, 2007) agrees that Stalnaker’s notion of a pragmatic presupposition should play a role in the semantics of ‘know’ when claiming that, within the framework of his contrastivist account, the contrasts relevant in *C* are ‘always recoverable’ from Stalnaker’s context set or that the context set ‘provides the default source of contrasts’ (Schaffer 2005, p. 249). However, Schaffer seems sceptical about the contextualist approach defended here when describing Lewis’s rules as ‘little more than a laundry list of rules of thumb, replete with unclear principles, subject to a variety of counterexamples, and open to skeptical usurpation as merely pragmatic’ ((Schaffer 2004a, p. 88), but see also (Schaffer 2005, p. 267)). More importantly, Schaffer explicitly rejects the idea of explicating Lewis’s notion of proper ignoring in terms of what is pragmatically presupposed. Schaffer: ‘if the contextualist deploys anything like Stalnaker’s notion of a context set, then [she] must forgo

such Lewisian Rules as Actuality, Belief, and Resemblance, since the context set need not contain actuality, need not correspond to anyone's beliefs, and is not closed under resemblance [. . .]. As such, contextualism would no longer underwrite, e.g. Lewis's solutions to skepticism, Gettier cases, and the lottery paradox, since these require Actuality, Belief, and Resemblance.' (Schaffer 2004a, pp. 99, fn. 27).

Considering my above formulation of (RP), however, it is fairly obvious that, *pace* Schaffer, the contextualist *can* deploy Stalnaker's notion of a context set in explicating the notion of proper ignoring. For further discussion of the interaction between (RP) and the remaining Lewisian rules see Ch. 5. For discussion of Schaffer's contrastivist account see (Kvanvig 2007; Neta 2008; Stalnaker 2004).

(23) What happens if your son refuses to pragmatically presuppose that you did not dream? In such a case you will find yourself in what Stalnaker (1978) calls a *defective context*. As I argue below, in defective contexts it is unclear whether you satisfy 'knows', this view providing an attractive explanation of our unclear intuitions about the acceptability of 'knowledge'-ascriptions in defective contexts (see Sect. 7, pp. 43-5).

(24) I again assume that none of the other Lewisian rules that (RP) is to be supplemented with prohibits properly ignoring sceptical possibilities in *C*.

(25) The importance of the idea that the conversational participants should have authority over the 'epistemic standards' of their own context has been emphasized by many contextualists in recent years. See especially (DeRose 2004b), but also (Cohen 1999; Neta 2002; and Schaffer 2005). As we shall see in greater detail below, authority over one's own epistemic standards can be made available by pairing (RP) with a suitable notion of pragmatic presupposition.

(26) Strictly speaking, Stalnaker gives a three-stage definition of the notion of pragmatic presupposition, the first step consisting of a definition of 'acceptance'. These details do not concern me here, however. I work instead with an intuitive notion of *acceptance for the purpose of one's conversation*.

See (Stalnaker 2002, p. 716) and (Stalnaker 1984, pp. 79–82) for the notion of acceptance.

(27) (Stalnaker 2002, p. 716) uses a simple conditional rather than a biconditional but considering that he aims to ‘define’ common ground, a biconditional appears more adequate here. Moreover, note that the relevant beliefs are implicit beliefs (see (Lycan 1986) for a discussion of implicit beliefs).

(28) (Stalnaker 2002, p. 707 and p. 717). Stalnaker has defended accounts of pragmatic presupposition similar to this one since at least (Stalnaker 1974, p. 49), while the general idea underlying the account can already be found in (Stalnaker 1970, pp. 38–40). Note also that I am not addressing issues arising from the topic of presupposition accommodation here. See (Stalnaker 2002, pp. 708–15) (esp. fn. 14) and (von Stechow 2008) for interesting discussion.

(29) See (Stalnaker 1978, pp. 89–90, 1998, p. 102, 2002, pp. 712–13), and also (von Stechow 2008) and (Yablo 2006, p. 165).

(30) Faculty Meeting is also a counterexample to the definition of pragmatic presupposition defended in (Soames 1982).

(31) (Stalnaker 1974, p. 52; emphasis added). A closely related passage is (Stalnaker 1978, p. 84): ‘A proposition is presupposed if the speaker is disposed to act as if he assumes or believes that the proposition is true, and as if he assumes or believes that his audience assumes or believes that it is true as well.’

(32) Note that a partial disposition of the relevant kind is, strictly speaking, sufficient for pragmatic presupposition. For instance, at the end of Faculty Meeting, A is disposed to assert ‘I have to pick up my sister from the airport’ but she is not disposed to answer ‘yes’ when asked whether it is common ground that she has a sister. Being disposed to answer ‘yes’ when asked whether *p* is common ground, however, is surely required for being *fully* (i.e. in all respects) disposed to behave, in one’s use of language, as if one believed *p* to be common ground. Thus, a full disposition of the relevant kind is not required for pragmatic presupposition and (PP*) is, strictly speaking, false: it needs to be qualified by inserting ‘partially’

into its right-hand side. In what follows I leave this qualification aside for stylistic reasons.

(33) It might be objected here that we have as much voluntary control over our attention as we have over our behavioural dispositions, because we typically can, when asked to attend to a particular object before us, freely decide to attend to it or not: we usually have, it seems, voluntary control over which objects we attend to. In response to this objection it is instructive to distinguish between the perceptual act of attending to physical objects (perception) and the intellectual act of attending to propositions or possibilities (thought). It is surely correct that perceptual attendance is subject to a large degree of voluntary control, but this does not seem to be the case with intellectual attendance, the notion at issue in Lewis's *Rule of Attention*: in order to decide whether one attends to a certain possibility, one needs to direct one's mind towards that very possibility and thus needs to attend to it. As a consequence, one cannot successfully decide not to attend to a certain possibility: acts of intellectual attendance are not subject to voluntary control in the way in which acts of perceptual attendance are. Moreover, note that when somebody mentions or expresses a possibility in conversation the listener attends to that possibility purely in virtue of cognitively processing and interpreting the speaker's assertions. In interpreting language, one inevitably directs one's mind towards the propositions and possibilities expressed by the speaker.

(34) A few remarks on the notion of direct voluntary control are in order. What is direct voluntary control? A state of affairs is under your direct voluntary control if your mere choosing to perform a certain action is sufficient to bring about that state of affairs. For instance, imagining that you have a red nose is, under normal circumstances, under your direct voluntary control, for as soon as you choose to imagine that you have a red nose, you imagine that you have a red nose. Similarly, your behavioural linguistic dispositions are, under normal circumstances, under your direct voluntary control: as soon as you choose to be disposed to assert sentences such as 'I have to pick up my sister from the

airport', you are disposed to assert sentences such as 'I have to pick up my sister from the airport'.

A given state of affairs is, however, under your indirect voluntary control if it is (a) not under your direct voluntary control, but (b) you can nevertheless bring about that state of affairs by choosing actions that bring it about. For instance, raising your blood pressure is under your indirect voluntary control: by merely choosing to raise your blood pressure, your blood pressure will not be raised. However, since you can choose to exercise in order to raise your blood pressure, you have indirect voluntary control over your blood pressure. Another example of indirect voluntary control is my current belief that there is a banana on my desk. I (presumably) cannot believe that there is a banana on my desk merely by choosing to believe that there is a banana on my desk (I should note that there is no banana on my desk), but I can choose to place a banana on my desk, which would bring about my believing that there is a banana on my desk. For further background on the distinction between direct and indirect voluntary control see (Alston 2005), Ch. 4.

(35) Of course, what a *participant to a conversation* takes seriously for herself can differ from what she takes seriously for the purposes of the conversation and thus from what is taken seriously in the conversation. I discuss the significance of this point in Sections 2.3–2.5.

(36) Note also that (PP*) does not define the notion of a pragmatic presupposition in terms of the *pretence* to believe that *p* is common ground, a strategy that has been criticized by Gauker (1998) and, building on Gauker's objections, by von Stechow (2008), who seems to reject (PP*) on the basis of Gauker's arguments. However, since pretending that one believes *p* and behaving linguistically as if one believed *p* are two entirely different notions, Gauker's arguments do not pose a threat to (PP*).

(37) It might also be useful at this point to remind us that both Lewis's account and PEC are effectively contextualized relevant alternatives accounts of 'knowledge'. Thus, the notion of what can be properly ignored at a context at work in the

above definitions should be understood as coinciding with the notion of what is *epistemically irrelevant* at a context: a world *w* can be properly ignored at a context *C* just in case it is epistemically irrelevant in *C*. Thus, when I say that the Rule of Presupposition places a constraint on the notion of proper ignoring, more traditional relevant alternatives theorists are invited to understand this as the rule placing a constraint on the notion of epistemic relevance at a context: it is that notion that, according to PEC, is to be partly understood in terms of pragmatic presuppositions.

(38) It should be noted that Lewis does not accept the view that 'knowledge' entails belief. Given Lewis's view that 'knowledge' does not entail belief, he would most certainly also reject the view that 'knowledge' entails the presence of a properly based belief. However, he would then still face the type of counterexample presented in the main text. To avoid those, he would have to at least add the constraint that the subject does not have an improperly based belief *p*. Thus, Lewis does not get around amending (L), if he takes seriously the type of example mentioned in the main text.

(39) Ichikawa (2011a, p. 386) also argues that a proper basing constraint should be added to Lewis's account, but his constraint demands that the subject's belief be based on her evidence, which is a stronger condition than the one explicated in the main text and to be developed in Section 5.3.

(40) The reader should already be warned that I shall later refine and amend some of the rules mentioned here.

(41) This problem has been brought to my attention by John Hawthorne and Nico Silins.

(42) See (Lewis 1986, Ch 1.4) and (Stalnaker 1984, 1987, 1988).

(43) For illustration, note that Lou's belief as to whether FLT is true is accidentally true only, for there are nearby worlds in which Lou comes to believe a random falsehood instead of FLT on the basis of tealeaves reading.

(44) One might think that we ought to merely attempt to give an account of empirical knowledge and address the issue of knowledge of necessary truths with a different, additional theory. But such a move would be problematic for two reasons. First, note that some knowledge of necessary truths is empirical knowledge. Simple examples can be construed by noting that any contingent truth can, by disjunction introduction, be transformed into a necessary truth. For instance, if Maya knows the contingent proposition p , but not the necessary proposition q , she can come to know the necessary proposition $(p \vee q)$ by competent deduction. Maya would, in the case imagined, know this disjunctive proposition at least partly on empirical grounds—namely, on the basis of her empirical evidence in support of p . Secondly, note that it would surely be desirable for purely theoretical reasons to offer an account that is not, in an ad hoc manner, restricted to contingent truths.

(45) Cp. (Williamson 2000).

(46) It is worthwhile emphasizing how very close this view is to (L*), if we think of it as explicating (L*)'s proper basing constraint in terms of safety and as rendering superfluous the Lewisian Rule of Actuality and Rule of Resemblance.

