

# Three Lewisian semantics for “know”

Julien Dutant, [julien.dutant@unige.ch](mailto:julien.dutant@unige.ch), comments welcome

15<sup>th</sup> june 2013

Epistemic contextualists hold that “know” expresses different relations between subject and fact in different contexts of use. They take the variation to explain why sceptical arguments are appealing but ultimately fail. In conversations in which we take sceptical scenarios seriously, “know” expresses a relation that almost no one has to any fact; but in ordinary conversations, it expresses a relation that many have to many facts. The sceptic’s denials of knowledge appear true because they are true, but they do not conflict with ordinary knowledge attributions.

Lewis (1996) has defended a substantial and independently motivated version of epistemic contextualism. It is substantial because it provides a fairly detailed semantics for “know”. It is independently motivated because it diagnoses other standing epistemological problems, namely the lottery paradox and the Gettier problem. The view has been influential, attracting both criticism and followers.<sup>1</sup>

---

<sup>1</sup>See [Cohen \(1998\)](#); [Hawthorne \(2002\)](#); [Schaffer \(2004\)](#) for criticism particularly targeted at Lewis’s semantics. Amended versions of the Lewisian semantics have been put forward by [Blome-Tillmann \(2009, 2012\)](#) and [Ichikawa \(2011a,b\)](#).

Yet for all the debate, a crucial aspect of Lewis's semantics has been left unspecified. As a result, a number of serious difficulties have been overlooked, as well as some ways to avoid them. We lay these issues out. We will not reach a final verdict on Lewisian contextualism but we will find all current versions wanting.

In a nutshell, the unspecified but crucial aspect is the following. In Lewisian semantics, context affects the value of "know" by selecting some possibilities as relevant. But there are three ways to do so. One is directly to select a set of possibilities as being relevant *simpliciter*. Another is to select a relation between cases and possibilities; a possibility is indirectly selected as relevant *to a case* if it bears that relation to that case. A third is to select some possibilities directly, some relationally. We call the respective semantics *rigid*, *relational* and *mixed*.

Lewis and his followers have not clarified which semantics they endorse. The letter of [Lewis \(1996\)](#) and [Blome-Tillmann \(2009\)](#) suggests a rigid semantics; charity enjoins us to attribute them a mixed one. [Ichikawa \(2011a\)](#) does not specify his proposal in enough detail to decide, but his analogy with quantifiers suggests a relational one.

Yet the choice matters. When we examine its consequences, we find a number of overlooked difficulties for Lewisian views. Rigid semantics is incompatible with the factivity of "know". Mixed semantics make implausible predictions and face a charge of philosophical irrelevance. In their most straightforward versions, relational semantics face a problem with speakers in sceptical scenarios. The difficulties are not lethal. Neo-Lewisians may bite some of the bullets or try to

develop an alternative relational semantics. Our task here is only to bring the issues into light.

We proceed as follows. Section 1 introduces Lewisian contextualism and its three variants. Section 2 deals with rigid semantics. It argues that they violate factivity (2.1). That may raise doubts as to whether any Lewisian ever endorsed the view; we provide some textual evidence that they did (2.2) and particularly stress the ambiguity of the “Rules” of Lewis’s semantics (2.3). Section 3 deals with mixed semantics. We point out a significant class of Lewisian accounts committed to mixed semantics (3.1) and argue that they have implausible consequences (3.2). More generally, we argue that mixed semantics face a charge of philosophical irrelevance (3.3). Section 4 deals with relational semantics. One proposal close to the spirit of Lewis’s runs into trouble because it relies on attention (4.1). Others run into trouble with cases involving speakers who are themselves in sceptical situations (4.3). A fully satisfactory option yet remains to be articulated (4.4).

## **1 Three types of Lewisian semantics for “know”**

### **1.1 Lewisian contextualism**

Lewis states his contextualist semantics for “know” as follows:

S knows proposition P iff S’s evidence eliminates every possibility in which not-P – Psst! – except for those possibilities that we are properly ignoring. (Lewis, 1996, 554)

Lewis is not meticulous on use and mention (1996, 566) and leaves time and world indexes implicit. He also appeals to specific notions of possibilities and propositions. Let us reformulate more rigorously. Our central notion will be that of a *case*: a subject at a certain time and world. Since cases specify a subject, world and time, and in order to avoid cumbersome indexing, we can treat ‘know  $p$ ’ as a predicate that applies to cases: where  $S$ ,  $t$  and  $w$  are the subject, time and world of a case  $c$ , we say that  $c$  *satisfies* ‘know  $p$ ’ iff  $S$  knows that  $p$  at  $t$ ,  $w$ .<sup>2</sup> A *proposition* is something that holds or fails to hold at each case. Some propositions are *de se*: for instance, the proposition *being presently hungry* holds at exactly all cases in which the subject of the case is hungry at the time and world of the case. A *possibility* is a set of cases (1996, 552). What Lewis calls *evidence* is the experiences and memories that the subject of the case has at the time and world of the case (1996, 553). He says that a case *eliminates* a possibility if and only if one’s evidence in that case differs from the evidence one has (in the cases of) that possibility (1996, 553). We can simply say that two cases are *internally equivalent* if the experiences and memories of their respective subjects at their respective times and worlds are the same. A case eliminates another just if they are not internally equivalent. We can now restate the first half of Lewis’s semantics:

A case  $c$  satisfies ‘knows  $p$ ’ iff  $p$  holds in all cases that are internally equivalent to  $c$ .

Equivalently: ‘ $S$  knows that  $p$ ’ is true at  $t, w$  iff for any  $S', t', w'$  such that  $S'$  has the same experiences and memories at  $t', w'$  as  $S$  has at  $t, w$ ,  $p$  is true at  $S', t', w'$ .

---

<sup>2</sup>We use single quotes as corner quotes.

So far the semantics has sceptical consequences, since for pretty much any case and proposition concerning external matters of fact there is an internally equivalent case in which the proposition is false. That is where the contextualist element comes in. On Lewis's view, in any conversation in which matters of knowledge are discussed, there is a range of possibilities that are "not properly ignored" by the speakers. We say that these possibilities are *relevant in that conversational context*. By extension, a case is relevant in a context if it belongs to a possibility that is relevant in that context.<sup>3</sup> Now Lewis says that in any given context, "know" expresses a relation that a subject bears to a proposition  $p$  whenever their evidence eliminates not- $p$  cases *that are relevant in that context*. Hence relative to a given context  $C$ , a case satisfies 'knows  $p$ ' if and only if  $p$  holds at all cases that are internally equivalent to  $c$  and *relevant in  $C$* .

---

<sup>3</sup>Let me flag a tricky issue here (see also [Hawthorne, 2002](#), 243). Suppose we are discussing whether Alice knows where her keys are. She believes that they are in the drawer. Taking Lewis's Rule of Belief ("a possibility that the subject believes to obtain is not properly ignored", [Lewis, 1996](#), 555) at face value, we should count the possibility *that the keys are in the drawer* as relevant. But Lewisians do not want all cases belonging to that possibility to be relevant. For they include cases in which the keys are in the drawer *and* Alice is in some sceptical scenario; Lewisians do not want those to be relevant just because Alice takes her keys to be in the drawer. There are two ways out of the issue. One is to say that the possibility Alice believes to obtain is more specific — that the keys are in the drawer and *everything else is normal*, for instance. Another is to say that only a *subset* of the cases belonging to relevant possibilities are relevant — the "most normal" ones, or some such idea. The problem is a symptom of a general tension between two uses Lewis makes of notion of possibility: *qua* things that the evidence eliminates, they need to be fairly specific; *qua* things people attend to or believe, they need to be fairly unspecific. Lewis brushes the issue aside by stipulating that "a possibility will be specific enough if it cannot be split into subcases in such a way that anything we have said about possibilities, or anything we are going to say before we are done, applies to some subcases and not to others" ([Lewis, 1996](#), 552). But as our example shows, it is unclear in some cases whether there is a level of specificity that allows him to say all what he wants to say about a certain possibility. We ignore the issue in what follows.

Lewis's semantics can be amended in some ways that are not crucial to our present discussion. As stated, it has highly problematic implications: logical omniscience (Lewis, 1996, 552), perfect knowledge of one's experiences and memory (Hawthorne, 2004, 60n) and irrelevance of bases of belief (Ichikawa, 2011a, 386; Blome-Tillmann, 2009, 259). In reply Lewis (*ibid.*) endorsed Stalnaker's defence of attitudes with coarse contents (Stalnaker, 1984, chap. 4). Neo-Lewisians have rather amended the semantics. In addition to having evidence that eliminates error, Blome-Tillmann (2009, 259) requires that one has a belief that is "properly based" and Ichikawa (2011a, 386) requires that one has a belief that is based on the evidence relevant for that belief.

So for our purposes, a "Lewisian" semantics will be one according to which a case satisfies 'knows  $p$ ' in a context just if  $p$  holds at all internally equivalent cases that are relevant in that context and, optionally, some basis requirement is satisfied.

## 1.2 Why Lewisians need internal equivalence

Here I must slightly digress. Ichikawa (2011a, 388) dismisses Lewis's internalist notion of evidence. On his brand of Lewisianism, elimination by evidence does not require internal distinctness. While that aspect of the semantics is mostly orthogonal to our discussion, we will at various points assume that sceptical scenarios cannot be eliminated. We cannot make that assumption unless elimination requires internal distinctness, so Ichikawa would object to applying those points to Lewisian views in general.

However, Lewisians cannot really afford giving up internal equivalence. We argue the point below. But it is mostly tangential to our discussion. Readers who are not convinced may take our targets to be “internalist” Lewisian views alone. Readers who are not interested in the issue can harmlessly jump to the next section.

Ichikawa suggests replacing Lewis’s internalist notion of evidence with a contextualist version of the idea that evidence is knowledge. Making use and mention explicit, the proposal is:

**Ichikawa’s non-reductive Lewisianism** A case  $c$  satisfies ‘knows  $p$ ’ relative to a context  $C$  iff there are some propositions  $e_1, \dots, e_n$  such that (a) in  $c$ , one believes that  $p$  on the basis of  $e_1, \dots, e_n$ , (b)  $e_1, \dots, e_n$  eliminate all relevant not- $p$  cases, and (c) *for each*  $e_1, \dots, e_n$ ,  $c$  satisfies ‘knows  $e_k$ ’ relative to  $C$ .

The condition for “know” to be satisfied is stated in terms of “know” being satisfied by the same case-proposition pair or another. As [Ichikawa \(2011a, 389\)](#) points out, that is not viciously circular. Rather, it puts constraints on how contexts may distribute instances of “know” in relation to facts about basing.

Still, the amendment seriously cripples the Lewisian explanation of sceptical paradoxes. On Lewis’s story, sceptical scenarios (internally alike error cases) cannot be eliminated. So bringing them into relevance is sufficient to falsify knowledge claims. On Ichikawa’s non-reductive proposal, when a sceptical scenario is brought into relevance, it is still an open question whether we eliminate it. For all the view says, even relative to a context in which a handless brain-in-a-vat sce-

nario is relevant, we may still satisfy “know that we have hands”. If we do, we eliminate the scenario in question.

So to explain how knowledge claims turn out false in sceptical contexts, Ichikawa not only needs to say (a) how the domain of relevant possibilities expands to include sceptical scenarios, but also (b) how the extension of “know” shrinks so that pretty much only internal facts count as “known” and thus “evidence” in those contexts. While he has quite a lot to say about (a), he says next to nothing about (b) (see [Ichikawa, 2011a](#), 390). And for the reasons just given, he cannot rely on (a) to explain (b). Yet it now looks as if (b) does most of the work. Once we have said how a sceptical context makes it so that no external facts count as “known” relative to them, there is not much left for uneliminated alternatives to do.

The problem with Ichikawa’s suggestion generalizes. The core of the Lewisian story is that bringing sceptical scenarios into relevance falsifies knowledge claims because those scenarios are error possibilities uneliminated by one’s evidence. Only internalist notions of evidence fit the story. On externalist notions, some sceptical scenarios would be eliminated. On a contextualist notion of evidence, the core of the story would be instead how focusing on sceptical scenarios makes us switch to internalist notions of evidence.

Lewisians need an internalist notion of evidence. More precisely, Lewisians need not follow Lewis in calling one’s internal state “evidence” — perhaps the term has better uses. But they do need internal distinctness (or something like it, *e.g.* internal dissimilarity) to play the role of what Lewis calls “elimination by one’s evidence”. Otherwise they cannot assume that sceptical scenarios are



“uneliminated”.

### 1.3 Three types of Lewisian semantics

When spelling out relevance, we face a choice. On one option, context provides a set of possibilities that are relevant *simpliciter*. Relative to a given context, any case is evaluated with respect to that very same set of possibilities. That is, a case  $c$  satisfies “knows  $p$ ” relative to a context  $C$  just if  $p$  holds in all cases that are internally equivalent to  $c$  and relevant (*simpliciter*) in  $C$ . We call that a *rigid* semantics. On another option, context provides *each case* with a set of possibilities that are relevant *to that case*. Relative to a given context, a case is evaluated with respect to the possibilities that are relevant *to that case*. That is, a case  $c$  satisfies ‘knows  $p$ ’ relative to a context  $C$  just if  $p$  holds in all cases that are internally equivalent to  $c$  and relevant *to  $c$*  in  $C$ . We call that a *non-rigid* semantics.<sup>4</sup>

To illustrate, suppose that some speakers discuss whether a man presently knows that his keys are in his backpack. We thus have a case,  $c$ , consisting of that man, the present time and the actual world; we also have a proposition,  $p$ , stating that the keys are in that man’s backpack. On Lewis’s semantics, whether  $c$  satisfies ‘knows  $p$ ’ by the standards of our context depends on what goes on in a number of relevant cases. To give a crude illustration, imagine that only two cases are relevant: the actual case and one in which the keys have been taken away.

---

<sup>4</sup>Some readers may find the distinction reminiscent of Cohen’s (1998, 294) distinction between “speaker-sensitive” and “subject-sensitive” selection of relevant possibilities. They are crucially different. See footnote 10 below.

Now suppose we ask whether some *other* case  $c'$  — say, a counterfactual case — satisfies ‘knows  $p$ ’ in the very same context. On a rigid semantics, that will again depend on what goes on in those particular two cases. On a non-rigid semantics, the cases relevant to whether  $c'$  satisfies ‘knows  $p$ ’ may differ from those relevant to whether  $c$  satisfies it.

We should further distinguish two types of non-rigid semantics. On a *relational* one, a context provides a relation between cases. A case is relevant to a given case  $c$  in that context if  $c$  is so related to that case. For a simple illustration, imagine that a context selects the identity relation. In that context, each case is the only case relevant to itself and “knows” denotes something like true belief. Alternatively, on a *mixed* semantics, some cases are selected relationally and some cases rigidly. For instance, relevant cases may be selected in virtue of being either (a) similar enough to the target case *or* (b) attended to by the speakers of the context. Cases picked up by the similarity relation will vary from one case to another; cases picked up by what speakers attend to will not.

To make our types exclusive and exhaustive, we define them as follows. If, for each context, there is a (non-empty) set of cases that is uniquely relevant to every case, the semantics is rigid. If, for each context, there is a set of cases that is relevant to every case but not uniquely relevant to some or all cases, the semantics is mixed. If, for some contexts, there is no set of cases that are relevant to every case, the semantics is relational.<sup>5</sup>

---

<sup>5</sup>We use extensional definitions so that one can easily check that they are exclusive and exhaustive. Their drawback is that they may misclassify some semantics that are mixed in spirit as relational and conversely. Consider a semantics such that for any context  $C$  and case  $c$ , a case  $c^*$

## 2 Rigid semantics

### 2.1 The factivity problem

With the options clearly set out, one sees that rigid semantics face a straightforward difficulty: they violate factivity.<sup>6</sup> Factivity is the idea that knowledge entails truth. In a contextualist setting, it is the idea that a case satisfies ‘knows  $p$ ’ relative to a context  $C$  only if it satisfies  $p$  relative to  $C$ .

Consider some ordinary context  $C$  in which speakers discuss an ordinary case  $c$  in which Alice sees a tree. Assume that the speakers truly utter (1):

(1) Alice knows that there are trees.

By the rigid semantics, there are trees at all cases that are internally equivalent to  $c$  and relevant in  $C$ . Now there is some possible world in which some subject, Ecila, has the same experiences and memories as Alice has in  $c$  but is the victim of a Cartesian demon. In particular, in Ecila’s case there are no trees. Call Ecila’s case  $c^*$ . By the semantics,  $c^*$  satisfies “knows that there are trees” if there are trees

---

is relevant to  $c$  in  $C$  iff (a)  $c^*$  is identical to  $c$  or (b)  $c^*$  is attended to in  $C$ . It should be classified as mixed. Yet under a liberal notion of context, there are contexts without participants. Since nothing is attended to in such contexts, condition (b) does not pick any case and no case is relevant *simpliciter*. Our definitions would wrongly count the semantics as relational. We can avoid that by implicitly restricting “each context” in our definitions to contexts in which knowledge is under discussion. Second, consider a semantics according to which  $c^*$  is relevant to  $c$  in  $C$  iff  $c$  is sufficiently similar to  $c^*$  by the standards of  $C$ . The semantics is relational in spirit. Yet it may in principle happen that some special case is sufficiently similar to any case by the standards of every possible context. If that was so, in each context, they would be a case that is relevant to any other — namely, the special case. Our definitions would wrongly count the semantics as mixed. While in principle possible, I doubt that any reasonable relational semantics would have that consequence. We can safely stick to our extensional definitions with the implicit restriction to contexts in which knowledge is under discussion.

<sup>6</sup>See also Stanley, 2005, 111–2.

at all cases that are internally equivalent to  $c^*$  and relevant in  $C$ . But since  $c^*$  and  $c$  are internally equivalent, any case internally equivalent to the first is internally equivalent to the second. So  $c^*$  satisfies “knows that there are trees” if there are trees at all cases that are internally equivalent to  $c$  and relevant in  $C$ . But we have seen that it is so. So  $c^*$  satisfies “knows that there are trees” without satisfying “there are trees”. Factivity fails.

The result is surprising, since Lewis’s semantics includes a Rule of Actuality that is meant to ensure factivity (Lewis, 1996, 554). In our terminology, the Rule says that the actual case is always relevant. As Lewis points out, it is ambiguous. If a speaker considers what a friend of theirs knew yesterday, there are two “actual cases” involved: the speaker’s and their friend’s. Lewis thus specifies the Rule: it only makes relevant *the subject’s* case (Lewis, 1996, 555). But that is still ambiguous. Which subject? There are many subjects. And which actuality of a subject? Each subject has many “actualities”.<sup>7</sup> Obviously, Lewis does not intend that for any context  $C$ , *each case in which the subject of the case exists at the time and world of the case* is relevant. Scepticism would readily follow. By “the subject’s actuality”, Lewis must mean one of two things: the case *under discussion in a given context*, or the case *of evaluation*. Hence two readings of the Rule:

**Rigid Rule of Actuality** For every context  $C$ , if  $c$  is a case under discussion in  $C$ , then  $c$  is relevant *simpliciter* in  $C$ .

**Relational Rule of Actuality** For every context  $C$  and case  $c$ ,  $c$  is relevant to  $c$  in

---

<sup>7</sup>Lewis takes subjects to exist at one world only; but as long as they exist more than one instant, even world-bound subjects are the centre of many cases.

*C*.

In a rigid semantics, one must adopt the rigid version. But the rigid Rule can only ensure factivity *for the cases under discussion*. In our example, it ensures that Alice's case *c* is relevant to *C*. So Alice satisfies 'knows *p*' in *c* relative to *C* only if *p* holds in *c*. But the rule fails to ensure factivity for *unmentioned cases*. Ecila's case *c\** is not relevant in *C*; hence whether it satisfies 'knows *p*' relative to *C* does not depend on what happens in it.

So the problem only affects unmentioned cases. Is it serious? Yes. First, truth-conditional semantics routinely assign truth-values to sentences (uttered or not) relative to circumstances of evaluation that are not under discussion. That is needed in order to give compositional accounts of quantified, tensed and modal sentences. As a result, the factivity problem crops up in uttered claims. Under some assumptions, "someone could know something false" comes out true relative to non-sceptical contexts.<sup>8</sup> Second, it is simply implausible that in most contexts *know* expresses a relation that some people bear to false propositions. That conflicts with much of what we think we know about knowledge: that it entails truth, that it is a worthwhile aim of inquiry, that it is a valuable state to be in, for

---

<sup>8</sup>The assumptions are: that "some" and "could" do not bring under discussion all the cases they range over, and that the range of "could" in a context typically extends beyond the cases that are relevant to "know" in that context. By the second assumption, some factivity-violating case typically enters the range of "some" and "could"; by the first, that does not bring those cases into discussion and hence relevance. One may dispute both assumptions for ordinary "could", but they clearly hold for "it is metaphysically possible that" or "Metaphysically speaking, ... could ...". The first is obvious. The second holds unless it is impossible to use the operator "It is metaphysically possible that ..." without landing oneself in a sceptical context. So the rigid Lewisian semantics must endorse true instances of "it is metaphysically possible that someone knows something false", which is bad enough.

instance. We would have to say that knowledge *often* or *sometimes* entails truth, that it is a worthwhile aim of inquiry *when true*, that knowledge *under discussion* is valuable, and so on. That seems absurd.

In short, either a rigid semantics makes all cases relevant *simpliciter* or it makes some cases not relevant *simpliciter*, hence not even relevant to themselves. The first option entails scepticism and the second violates factivity. To avoid that, the *relational* Rule of Actuality is needed, which ensures that relative to any context, each case is relevant *to itself*, whether or not it is mentioned. Hence relative to any context, a case satisfies ‘knows *p*’ only if *p* holds in it. Factivity is preserved.

An analogous problem arises with Lewis’s Rule of Resemblance. The Rule is meant to deal with Gettier cases. But when fitted in an rigid semantics, it only ensures that Gettier-style situations *under discussion* fail to satisfy “know”. A relational version of the rule avoids the problem.

## 2.2 Does anyone endorse a rigid semantics?

Rigid Lewisian contextualism is untenable. In reply to that, some readers have doubted whether anyone has ever endorsed the view. In this they are moved by charity rather than textual evidence.

*Prima facie*, Lewis appears to endorse a rigid semantics. He writes:

“*S* knows that *P* iff *S*’s evidence eliminates every possibility in which not-*P* — Psst! — except for those possibilities that we [speakers] are properly ignoring” (Lewis, 1996, 554, see also 561).

He does not write: “. . . except for those possibilities that we are properly ignoring *relative to S and the time and world of S in question*”. But the qualification would be needed on a non-rigid semantics. Again, when he restates his semantics in terms of “proper presupposition”, he writes:

“*S* knows that *P* iff *S*’s evidence eliminates every possibility in which not-*P* except for those possibilities that conflict with our proper presuppositions” (*ibid.*, 554).

He does not write: “the possibilities that conflict with our proper presuppositions *relative to S and the time and world of S in question*”. In fact, relativizing presuppositions in that way may not even make sense — see section 3.1. Yet that would be needed on a non-rigid semantics.

Accordingly, commentators describe Lewis’s semantics as if it was rigid. Here is Ichikawa, for instance:

On Lewis’s view, there is a class of possibilities that grows and shrinks according to conversational context; our knowledge attributions are true when the subject’s evidence eliminates all of the members of that class in which the object of knowledge is false. (Ichikawa, 2011a, 385)

Ichikawa talks of a single class of possibilities that varies with context and gives no hint that there would be one such class per subject, time and world. He is not alone in this. None of the commentators I am aware of explicitly ascribes Lewis

a non-rigid semantics. At best, their description is ambiguous between the rigid and non-rigid ones.<sup>9</sup>

Blome-Tillmann (2009) and Ichikawa (2011a,b) defend amended versions of Lewis's semantics. Their statements are equally ambiguous. Blome-Tillmann (2009, 245) endorses the general form of Lewis's semantics, which he states as follows:

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

Again, there is no indication that the worlds are properly ignored in  $C$  relative to  $x$  and the time and world in question, as a non-rigid semantics requires. All of Lewis's Rules are endorsed except the Rule of Attention, which he replaces by a Rule of Presupposition. The latter is rigid — see section 3.1. So like Lewis's, Blome-Tillmann's statement suggests a rigid semantics but allows for a mixed one.

Ichikawa (2011a, 387) takes Lewis's Rules to be at best approximately correct. He rather develops Lewis's (1996, 553) passing analogy between "know" and quantifiers: relative to a context, one counts as "knowing" if one eliminates the alternatives that count as "all alternatives" in that context. His statement of the semantics leaves relativity to contexts, time and worlds implicit:

S knows that  $p$  just in case, for some evidence  $E$ , (i) S believes that

---

<sup>9</sup>See (Cohen, 1998, 290–1), Vogel (1999, 158), (Pryor, 2001, 97–8), (Feldman, 2001, 66), Schaffer (2001, 201; 2005, 125–6), Hawthorne (2002, 242; 2004, 60), Bach (2005, 81–2), Douven (2005, 573), Blome-Tillmann (2009, 245), Ichikawa (2011b, 295).



p on the basis of E, and (ii) all the E cases are p cases. (Ichikawa, 2011a, 387, see also 2011a, 295)

Where “‘knowledge’ inherits its context-sensitivity from the context-sensitive ‘every possibility’” (2011a, 385). Now there are two views on how context provides quantifiers with a restricted domain, a rigid and relational one. On the former, context rigidly provides a set of individuals. On the latter, it provides a relation between circumstances of evaluation and sets of individuals. Insofar as the relational semantics is preferred for quantifiers (Stanley and Szabó, 2000, 252), it may be natural to read Ichikawa that way. However, as we noted, his statement gives no hint of non-rigidity.

### 2.3 The ambiguity of Lewis’s Rules

Lewis does provide a set of Rules that jointly determine which cases are relevant in a given context. One would expect them to settle the matter; but surprisingly, they do not. One Rule is clearly rigid, but the others have rigid and non-rigid readings. For all the Rules say, Lewis’s semantics could be either rigid or mixed.

Some classifications will be useful. A Rule is *rigid* if it picks a set of cases; it is *relational* if it picks a relation. We treat any mixed Rule as a conjunction of rigid and relational Rules. A Rule is *context-sensitive* if what it picks up varies with context; it is *invariant* otherwise. A context-sensitive Rule is *topic-sensitive* if what it picks up is a function of case(s) under discussion; it is *speaker-sensitive* otherwise. Invariant Rules have the same outcome in all contexts; topic-sensitive

Rules have the same outcome in any two contexts that discuss the same case(s); speaker-sensitive Rules may yield different outcomes in any two contexts.<sup>10</sup>

Lewis introduces eight Rules. Four bring cases into relevance — Actuality, Belief, Resemblance and Attention — and four defeasibly exclude cases from relevance — Reliability, two rules of Method and Conservatism. The latter are defeasible because the former always take precedence.

The Rule of Attention is indisputably rigid:

**Rule of Attention** [...] a possibility not ignored at all is *ipso facto* not properly ignored. What is and what is not being ignored is a feature of the particular conversational context. (Lewis, 1996, 559)

People attend to possibilities *simpliciter*, not relatively. We do not attend to the possibility that Alice faces an artificial tree *relatively* to the case in which she faces a tree, for instance. We simply attend to it — or not. Hence the Rule picks up cases

---

<sup>10</sup>Cohen (1999, 293–4) distinguishes “speaker-sensitive” and “subject-sensitive” Lewisian rules. The latter are those that determine relevance in virtue of facts about “the subject”. Since “the subject” is ambiguous between case under discussion and case of evaluation, they include both invariant relational Rules and topic-sensitive (rigid or relational) Rules. Topic-sensitivity is also distinct from the “subject-sensitivity” in “subject-sensitive” (or “interest-relative”) invariantism (Hawthorne, 2004; Stanley, 2005). On the latter kind of view, whether one knows is sensitive to one’s practical interests. In a Lewisian framework, that requires a relational Rule according to which  $c^*$  is relevant to a target case  $c$  if  $c^*$  is, so to speak, *too close for one’s stakes in  $c$  to  $c$* . That contrasts with a topic-sensitive Rule according to which the stakes *of the subject under discussion* determine relevance. Let  $c^h$  be a high-stakes case and  $c^l$  a low-stakes one. Let  $C^h$  be a context in which  $c^h$  is under discussion and  $C^l$  a case under which  $c^l$  is under discussion. The invariant interest-relative Rule ignores the difference between  $C^h$  and  $C^l$ : in both contexts, the range of cases relevant to  $c^h$  will *ceteris paribus* be wider than the range of cases relevant to  $c^l$ . By contrast, the contextualist topic-sensitive Rule ignores the difference between  $c^h$  and  $c^l$ : in  $C^h$ , both get wide ranges of relevant cases, in  $C^l$ , both get narrower ones. Similarly for a speaker-sensitive Rule that is sensitive to speaker’s stakes. The three forms of interest-relativity can in principle be combined: the range of cases relevant to a case  $c$  in a context  $C$  would be sensitive to the stakes in  $c$ , the stakes of the speakers of  $C$  and the stakes of whomever they talk about.

that are relevant *simpliciter*, not relatively to other cases. It is speaker-sensitive and rigid.<sup>11</sup>

The Rule of Conservatism is also most naturally read as rigid (Lewis, 1996, 559). It roughly says if some possibilities are commonly ignored by people around “us”, then they may be ignored. In Lewis’s paper “us” may refer to any of: his readers, speakers of some context, subjects under discussion or subjects of evaluation. Here, however, the text makes clear that he intends those who ignore possibilities, *i.e.* the speakers. So we seem to have a speaker-sensitive rigid rule, namely:

**Rigid Rule of Conservatism** For each context  $C$ , if people around the speakers of  $C$  commonly know themselves to normally ignore some possibility that includes a case  $c$ , then, defeasibly,  $c$  is not relevant in  $C$ .

With some stretch, however, one could read Lewis as suggesting that speakers may ignore the *kinds* of possibilities that are commonly ignored. That could receive a relational reading:

**Relational Rule of Conservatism** For each context  $C$  and case  $c$ , if people around the speakers of  $C$  commonly know themselves not to count as relevant to a case cases that are  $R$ -related to it, then, defeasibly, cases that are  $R$ -related to  $c$  are not relevant to  $c$  in  $C$ .

---

<sup>11</sup>People may also attend to *relations* or *respects of resemblance* themselves. That leads to a relational version of the Rule of Attention, to which we return section 4.1. But that does not invalidate the present point. Even when we attend to a relation, we attend to it *simpliciter*, not relatively to something else. Since Lewis’s Rule of Attention is framed in terms of attention to *possibilities*, it is not relational.

So the Rule of Conservatism is not wholly unambiguous.

The Rule of Actuality is ambiguous between a topic-sensitive rigid one and an invariant relational one, as we have seen. The same ambiguity is found in the Rule of Belief and the two Rules of Method (Lewis, 1996, 555, 558). The Rule of Belief says that a possibility that “the subject” believes to obtain is relevant; the first Rule of Method says that a possibility in which “the subject”’s sample is not representative may be ignored; the second Rule of Method says that a possibility in which the best explanation of “the subject’s” evidence is false may be ignored. As before, “the subject” may refer to the case under discussion or cases of evaluation:

**Rigid Rule of Belief** For every context  $C$  in which a case  $c$  is under discussion, if  $c^*$  is compatible with what the subject believes in  $c$ , then  $c^*$  is relevant *simpliciter* in  $C$ .

**Relational Rule of Belief** For each context  $C$  and case  $c$ , if  $c^*$  is compatible with what the subject believes in  $c$  then  $c^*$  is relevant to  $c$  in  $C$ .

**Rigid Rules of Method** For every context  $C$  in which a case  $c$  is under discussion, if  $c^*$  is a case in which (a) the sample the subject has in  $c$  is not representative, or (b) the best explanation of the evidence the subject has in  $c$  is false, then  $c^*$  is defeasibly not relevant in  $C$ .

**Relational Rules of Method** For every context  $C$  and case  $c$ , if  $c^*$  is a case in which (a) the sample the subject has in  $c$  is not representative, or (b) the best

explanation of the evidence the subject has in  $c$  is false, then  $c^*$  is defeasibly not relevant to  $c$  in  $C$ .

The rigid rules are topic-sensitive, the relational ones invariant.

The Rule of Reliability is four-ways ambiguous. It is introduced as follows:

**Rule of Reliability** Consider processes whereby information is transmitted to us: perception, memory, and testimony. These processes are fairly reliable. Within limits, we are entitled to take them for granted. We may properly presuppose that they work without a glitch in the case under consideration. Defeasibly - very defeasibly! - a possibility in which they fail may properly be ignored. (Lewis, 1996, 558)

On a first reading, “they” rigidly refers to *perception, memory and testimony*. That is, the Rule is:

**Absolute Rule of Reliability** For every context  $C$ , a case  $c$  in which perception, memory or testimony fail is defeasibly not relevant in  $C$ .

That Rule is rigid and invariant. On a second reading, “they” refers to *reliable processes* and it is contingent which processes are reliable.<sup>12</sup> Since reliability is contingent, “reliable” should be relativized to worlds, or better, cases.<sup>13</sup> But

---

<sup>12</sup>A non-contingent notion of reliability could be devised. For instance, reliability could mean reliability *at normal worlds*, where we assume that there is some absolute notion of normal worlds. (Whether the assumption is defensible is another matter.) But if reliability is non-contingent, then whichever processes are actually reliable are necessarily so. That takes us back to an absolute version of the Rule.

<sup>13</sup>Relativization to cases allows processes that are reliable on a planet but not on another, in a century but not in another, and so on. If one does not think reliability varies in such a way, one can still relativize reliability to cases but take it to be invariant across co-worldly cases.

which case? The case of the speaker, the case under discussion or the case under evaluation? The three choices are open:

**Speaker-sensitive Rule of Reliability** For every context  $C$ , if  $c$  is a case in which a process that is reliable *in*  $C$  fails,  $c$  is defeasibly not relevant in  $C$ .

**Topic-sensitive Rule of Reliability** For every context  $C$  in which a case  $c$  is under discussion, if  $c^*$  is a case in which a process that is reliable *in*  $c$  fails,  $c^*$  is defeasibly not relevant in  $C$ .

**Relational Rule of Reliability** For every context  $C$  and case  $c$ , if  $c^*$  is a case in which a process that is reliable *in*  $c$  fails,  $c^*$  is defeasibly not relevant *to*  $c$  in  $C$ .

The first two versions readings are context-sensitive and rigid, the third invariant and relational.

Consider finally the Rule of Resemblance:

**Rule of Resemblance** If one possibility saliently resembles another and the latter is not properly ignored (in virtue of rules other than this rule), then the former is not properly ignored. (cf. [Lewis, 1996, 556](#)).

The Rule is parasitic on other Rules. If all other Rules are rigid, it is rigid, otherwise it is relational.

**Rigid Rule of Resemblance** For each context  $C$ , if  $c$  is relevant *simpliciter* in  $C$  (in virtue of rules other than Resemblance) and  $c^*$  saliently resembles  $c$ , then  $c^*$  is relevant in  $C$ .

**Relational Rule of Resemblance** For each context  $C$  and case  $c$ , if  $c^*$  is relevant to  $c$  in  $C$  (in virtue of other rules than Resemblance) and  $c^{**}$  saliently resembles  $c^*$ , then  $c^{**}$  is relevant to  $c$  in  $C$ .

As Cohen (1998, 294–5) stresses, Lewis’s requirement that resemblance be *salient* is a further source of ambiguity. Salient to whom? Speaker, subject under discussion, or subject of evaluation? Lewis uses shifts in salience to explain shifty intuitions about lottery cases (Lewis, 1996, 557, 565–6). But as Cohen points out, the application involves a self-ascription in which speakers and subjects (both of discussion and evaluation) are the the same people. So it fails to disambiguate. Cohen claims that the Rule is disambiguated by its other intended application, namely to Gettier-style cases (1998, 295). On Lewis’s view, a Gettier-style case fails to satisfy “know” because it saliently resembles an error case. Cohen argues that it cannot be salient *to the subject in the Gettier case* that their case resembles an error case; so Lewis must intend the speakers. But as Ichikawa (2011a, 394–5) points out, Cohen assumes that two cases *saliently* resemble each other it is salient *that they resemble each other*. That need not be so. On a more liberal view of salient resemblance, two cases may saliently resemble each other because their *respects of resemblance* are salient: they share a property that is salient, or they stand in a relation that is salient. So the application to Gettier cases fails to disambiguate.

In conclusion, Lewis and his followers have been less than fully explicit in stating their semantics. On the one hand their formulations give no hint of non-rigidity. The could be easily interpreted as rigid. On the other hand, consideration

of unmentioned cases shows rigid semantics to be hopeless. Moreover, Lewis was well aware of the relational (hence non-rigid) nature of the standard semantics for modals and of Hintikka's (1962) semantics for "know", both of which have plausibly inspired his own semantics for "know".<sup>14</sup> So charity may lead one to think that Lewis and his followers had in mind a non-rigid semantics all along, though they did not feel the need to make it explicit.

We need not settle the matter here. Our aims so far were only to draw attention to the distinction between rigid and non-rigid semantics, to establish that Lewisians need a non-rigid semantics, and to argue that they need to be more explicit on that score.

### **3 Mixed semantics**

To avoid the factivity problem, Lewisians are committed to a relational Rule of Actuality, that is, to the idea that each case is relevant to itself, relative to any context. But to diagnose scepticism, many Lewisians need a rigid Rule. Being committed to both rigid and relational Rules, these Lewisians are committed to a mixed semantics. In the first section, I delineate a class of Lewisian who do need a rigid Rule. In the next two sections, we highlight two important difficulties that they face.

---

<sup>14</sup>See in particular the parallel treatment of modals and "know" in Lewis (1979, 355).



### 3.1 Taking possibilities into consideration and relevance *simpliciter*

Lewisians are committed to mixed semantics if they fulfil two conditions: first, to endorse the idea that sceptical arguments exert their force by making speakers *take certain possibilities into consideration*; second, to implement the idea in the most straightforward way: the possibilities in questions are made relevant *simpliciter*. The two conditions are sufficient (though not at all necessary) for Lewisians to be committed to a mixed semantics. Let us spell them out.

First, many Lewisians are committed to the idea that sceptical arguments exert their force by making speakers take certain possibilities into consideration — in some relevant sense of “taking into consideration”. Consider Lewis first. On his view, sceptical arguments call our *attention* to sceptical scenarios, that is, internally alike error possibilities. By doing so, they generate a context in which these possibilities are relevant. But by design, these possibilities cannot be eliminated. Hence most knowledge ascriptions turn out false in that context.

Both critics and followers of Lewis have objected to his Rule of Attention. *Merely attending* to sceptical scenarios does not seem sufficient to trigger sceptical judgements.<sup>15</sup> For instance, we may take some utterances of (2) to be true:

- (2) I know that many philosophers are not worried by Descartes’ Evil Demon scenario.

Yet such utterances draw attention to sceptical scenarios. If Lewis’s Rule applied,

---

<sup>15</sup>Oakley, 2001, 325–6; Williams, 2001, 15; Hawthorne, 2004, 64; Blome-Tillmann, 2009, 246–7; Ichikawa, 2011a, 388

the utterance would be false. If, as contextualists assume, our judgements track the truth-value of knowledge ascriptions, the rule is too strong.

Neo-Lewisians standardly revise Lewis's account as follows: possibilities are brought into relevance when speakers *take them seriously*, where merely attending to a possibility is not enough to take it seriously.<sup>16</sup> Blome-Tillmann (2009) has given the most detailed proposal along these lines. He proposes to cash out the relevant notion of taking a possibility seriously in terms of pragmatic presuppositions. He thus replaces Lewis's Rule of Attention with a Rule of Presupposition: if a possibility is compatible with what speakers presuppose in a context *C*, then that possibility is relevant to *C*. In contexts in which sceptical scenarios are presupposed not to obtain, knowledge attributions may come out true; in contexts in which sceptical scenarios are not presupposed not to obtain, they typically come out false.

Despite their differences, Lewis's Attention account and the various Taking Seriously accounts are variants of the idea that sceptical arguments exert their force by making speakers take certain possibilities into consideration. Attending to a possibility is a weak way of taking it into consideration; taking it seriously a more demanding one.

Now, the various attitudes of taking possibilities into consideration deliver a *rigid* parameter. As we noted, we attend to possibilities *simpliciter*, not relatively to cases — see 2.3. Similarly, on any reasonable notion of presupposition, what we presuppose are *possibilities*. One may presuppose *that one is seated by the fire*,

---

<sup>16</sup>Blome-Tillmann, 2009, 247–8, Ichikawa, 2011b, 296; see also Hawthorne, 2004, 64.

for instance. Possibilities determine sets of cases: cases compatible with them and cases incompatible with them. They do not determine relations between cases. It would make no sense to talk of a *relation between cases* being compatible with one's presuppositions or of presupposing something *relatively to a case*. Hence presuppositions only deliver rigid parameters. Similarly, on any notion of taking seriously that is relevant here, what we take seriously are possibilities. It would make no sense to talk of someone taking seriously the possibility that one dreams *relatively to a case in which one is awake*, for instance, or of someone talking seriously *the relation between cases of occurring in spatio-temporally close areas*. Thus what we take seriously must yield a rigid parameter.

Second, the most straightforward way to implement the Taking Possibilities into Consideration diagnosis of scepticism is through a rigid Rule. As we just argued, the possibilities taken into consideration in a given context determine a rigid set of cases. Now we simply add that those cases are relevant *simpliciter* in that context. The resulting Rule is rigid. There are more indirect ways to implement the diagnosis, which are compatible with a relational semantics — see section 4.2. But the most straightforward one is certainly through a rigid Rule.

So a significant class of Lewisian accounts are committed to a mixed semantics. These are accounts that rely on some variant of the idea that sceptical arguments exert their force by drawing sceptical scenarios into consideration, and that implement the idea in a Rigid rule according to which the scenarios taken into consideration in a context are relevant *simpliciter* in that context. Lewis's and Blome-Tillmann's accounts are cases in point. But more generally, since factivity

requires a relational rule, Lewisians are also committed to mixed semantics as soon as they endorse some rigid Rule.

### **3.2 The implausible asymmetries problem**

We now lay out two difficulties for mixed semantics. The first concerns only those who endorse the (rigid version of the) Taking Possibilities Into Consideration diagnosis of sceptical arguments. That is an important class of Lewisian views, as we stressed. The difficulty is that they predict implausible truth-value asymmetries.

To take a familiar example, suppose that a young Alice visits a zoo. She has seen the wolves and is now looking at the zebras. Two bystanders are taking into consideration (in whatever sense is relevant) the possibility that the zebras are cleverly painted mules. Noting that Alice's evidence cannot eliminate that possibility, they assert:

- (3) Alice does not know that the animals in that pen are not cleverly disguised mules.

By Lewisian lights, that sentence is true in their context. But now consider the following sentence:

- (4) Alice does not know that the animals in the previous pen are not cleverly disguised dogs.

Crucially, the speakers do not utter the latter sentence. They do not take into consideration the possibility that the animal labelled as wolves are cleverly disguised

dogs. Others things being equal — assuming in particular that the possibility is not actual or very similar to actuality —, the possibility is not relevant in their context. By Lewisian lights, sentence (4) is false in their context.

The asymmetry is implausible. If Alice does not count as “knowing” that the zebras are not cleverly disguised mules, she should not count as “knowing” that the wolves are not cleverly disguised dogs either. In both cases all she relies on is her own perception and her trust in her companions and the zoo. As for her perception, it is more easily fooled by disguised dogs than disguised mules — to untrained eyes, well-chosen dogs look more like wolves than mules look like zebras. As for her companions, they are more easily fooled by disguised dogs, for the same reasons. As for zoo authorities, there need not be any drastic asymmetry between their liability to use fake zebras and their liability to use fake wolves. The first has perhaps a greater payoff but the latter is less noticeable. In a nutshell, it is not in any sense harder for Alice to know that the zebras are zebras than that the wolves are wolves. Yet the semantics allows that (3) is true but (4) is false in the context of those speakers.

To avoid the consequence, Neo-Lewisians might say that the fake wolves possibility *is* taken into consideration, after all. For instance, they may argue that if speakers take into consideration the possibility that the zebras are fake, they distrust the zoo authorities enough so much that, implicitly or explicitly, they also take into consideration the possibility that the wolves are fake. More generally, they may assume that possibilities taken into consideration obey some resemblance constraint. For instance, they may say that some possibilities are primarily

or explicitly taken into consideration, and that any possibility which sufficiently resembles those (by some standard) also counts as taken into consideration. But it is unclear whether any *bona fide* notion of “taking into consideration” delivers that result. Lewis’s attention does not. Speakers may undoubtedly attend to the fake zebra possibility without attending to the fake wolves one. Blome-Tillmann’s presupposition does not either. Speakers may fail to presuppose that the animals in the present pen are zebras without failing to presuppose that the wolves are wolves. They may be wolves experts, for instance, who have no doubt whatsoever about the wolves but are less confident about the zebras. Thus they may perfectly presuppose that the wolves are not fake without presupposing that the zebras are. The point generalizes to any notion of “taking into consideration” under which the speakers do not take a possibility into consideration if they take themselves to know that it does not obtain.

Alternatively, Neo-Lewisians may endorse the consequence. They may argue that the truth value of knowledge ascriptions reflects the speakers’ own epistemic position, interests and worries as well as the subject’s position. Hence a subject may be equally well positioned with respect to  $p$  and to  $q$  and yet count as “knowing”  $p$  but not  $q$  relative to a context in which speakers are particularly worried about the latter. While that line may be taken, it fuels the worry that ordinary uses of *know* are philosophically irrelevant, to which we turn in the next section.

In reply to a somewhat similar case, [Blome-Tillmann \(2012\)](#) has revised his semantics in order to allow possibilities to be taken into consideration even when speakers take themselves to know that they do not obtain. The revision may be

thought to avoid the problem. But it does not. The proposed revision is an additional Rule of Evidence-Based Ignoring:

**Rule of Evidence-Based Ignoring** For any context  $C$ , case  $c$ , if the speakers in  $C$  ignore  $c$  because  $c$  is eliminated by their evidence, then  $c$  is not relevant *simpliciter* in  $C$ .

The Rule prevents the implausible asymmetry in our case provided that the speakers are ignoring the fake wolves possibility *because it is eliminated by their evidence*. That would be so, for instance, if the speakers are ignoring that possibility because they are wolves experts. But that would not be so if the speakers are ignoring it simply because it did not cross their mind. So the amendment fails to deal with all versions of our problem.<sup>17</sup>

---

<sup>17</sup>There are further worries with Blome-Tillmann's revision. First, it far from obvious that presuppositions can be sorted out into those that are made because our evidence entails them and others. For instance, you may presently presuppose that you are not a brain in a vat undergoing the experience of flying in the sky. Do you do so because your current experience eliminates that possibility, or because the possibility just did not cross your mind? Second, Blome-Tillmann does not discuss cases where we mistakenly take some possibility to be eliminated by our evidence. Nor does he discuss cases where a possibility is so eliminated yet not presupposed to obtain. It is not clear whether the Rule would give plausible results in such cases. Third, the account threatens to give up Blome-Tillmann's initial hope of accounting for "taking seriously" in terms of presupposition (2009, 247). Consider the following picture: there is a set of possibilities speakers take seriously; some of them are (taken to be) eliminated by their evidence, and they presuppose them not to obtain; others are not presupposed not to obtain; but all those that are taken seriously are relevant for their knowledge attributions. On that picture the notion of taking seriously is primitive and not explained in terms of presupposition; and it is doing all the work: what is presupposed or not is irrelevant. Blome-Tillmann's revised account is dangerously close to that picture. The only difference is the apparent reduction of taking seriously to presupposition along the following lines: a possibility is taken seriously by the speakers iff it is either compatible with their presuppositions or such that they presuppose it not to obtain only because their evidence eliminates it. Or alternatively: if and only if it is either compatible with their presuppositions or such that had their evidence not eliminated it, they would not have presupposed it to obtain (see the counterfactual version of the Rule in Blome-Tillmann, 2012, 117). Given the poor track record of counterfactual reductions, there is some reason to be pessimist about the second option. As to the

To sum up, on the most natural notions of taking possibilities into consideration, speakers do not take into consideration possibilities that they either do not think of at all or take themselves to eliminate. As a result, mixed semantics who implement the Taking Possibilities Into Consideration diagnosis in a Rigid rule predict implausible asymmetries in the truth-value of knowledge ascriptions. The asymmetries are a cost to the views.

### **3.3 The philosophical irrelevance problem**

The second worry for mixed semantics is more general. It is not directly an objection to the idea that the mixed semantics is the correct semantics for “know”. Rather, it is the worry that if ordinary uses of “know” have a mixed semantics, then the semantics of the ordinary language term “know” is of little relevance to epistemology.

Sosa (2000, 1–10) has argued that contextualism about “know” (of any type) is mostly irrelevant to epistemology. On his view, epistemologists ask whether we know anything and whether people speak truly when, in ordinary contexts, they claim to know things. Sosa grants that if epistemic contextualists are right, people often speak truly when they say, “I know such and such”. But it does not follow that we know anything nor even that people speak truly when, in ordinary contexts, they claim to know things. For by the contextualist’s own lights, that does not follow unless whatever satisfies “know” in their context satisfies it in our

---

first, it seems to us that, if anything, reduction goes the other way round. What makes it so that a possibility is presupposed not to obtain “only because it is eliminated” is that it is taken seriously, but seen to be eliminated.



present context. So, concludes Sosa, epistemic contextualism does not answer the question whether we know anything and it is of limited relevance to epistemology (Sosa, 2000, 3–4).

Sosa’s worry is not entirely clear. Blome-Tillmann has successfully answered some ways of construing it (Blome-Tillmann, 2007; Blome-Tillmann, 2009, 285–7). On one reading, contextualists do justice to our anti-sceptical intuitions only if they can claim that *people know things* and not merely that some utterances of “people know things” are true. On a second reading, epistemic contextualists are not doing epistemology as long as they merely mention “know” instead of using it. To the first objection, Blome-Tillmann replies that the anti-sceptical intuition is only that some utterances of “people know such and such” are true. To both objections, he replies that on his presupposition account, contexts of epistemological discussion need not be sceptical and hence allow for contextualists to claim that *people know such and such*. We may add a further reply. An epistemologist asks whether people know things. If “people know things” is true in *the epistemologist’s* context, then the answer to her question is positive. So insofar as the contextualist semantics applies to the epistemologist’s context, contextualism is relevant to epistemology.

Now the present worry is precisely this: if the mixed semantics is correct, it is doubtful that *epistemologists* should use “know” with its mixed semantics. If they should not, the contextualist’s mixed semantics is mostly irrelevant to epistemology. Thus a version of Sosa’s irrelevance worry has not been addressed. Let us detail how it arises.

Some semantic values are less natural than others. For instance, the semantic value of *being an umbrella or a table* is less natural than that of *being an umbrella*. On the mixed semantics, the semantic values that *know* receives in most contexts is unnatural in the way some conjunctive predicates are. Consider first an analogy. We could introduce a predicate *mome* such that '*x* is mome' is true in a context *C* iff *x* is near (a) its home and (b) the speaker of *C*. Part (a) of the semantics is relational and invariant. It selects a location relatively to the object to which the predicate is applied. Part (b) is rigid and context-sensitive. It selects a location rigidly as a function of the context. The predicate is equivalent to the conjunction of *being home* and *being near me*. Now consider a template mixed semantics such as the following:

**Template mixed semantics** For every context *C* and case *c*, *c* satisfies 'knows *p*' relative to *C* iff *p* holds in all cases internally equivalent to *c* that are (a) sufficiently similar to *c* and (b) taken into consideration by the speakers of *C*.

On such a semantics, '*knows p*' is equivalent to the conjunction of *has evidence that eliminates not-p cases that sufficiently resemble their case* and *has evidence that eliminates not-p cases that I take into consideration*. The first conjunct is relational and invariant; the second rigid and context-sensitive.

It is difficult to say what naturalness amount to in general. But we can make intuitive judgements about it and we can rely on some rules of thumb. In the case of *mome*, it is intuitively clear that the predicate is less natural than each of the conjuncts jointly equivalent to it. In the case of the mixed semantics, it is

not clear that the predicate is less natural than the second conjunct — the second conjunct violates factivity, for instance. But it is intuitively clear that it is less natural than the first conjunct. Furthermore, as a general rule, a conjunction of a relational condition and a rigid condition is less natural than a conjunct condition alone. The predicate *mome* is a case in point. Similarly, a mixed semantics will *ceteris paribus* be less natural than a rigid or a relational one. Finally, as we saw in the previous section, the kind of mixed semantics Lewisians typically endorse results in asymmetries in the truth-value of knowledge ascriptions that do not reflect natural differences in the epistemic positions of subjects but rather a combination of the latter with the speakers' own epistemic position, interest and worries. Together, these considerations strongly suggest that the mixed semantics is less natural than the one we get by dropping the rigid conjunct.

Lewis himself was well aware that his account assigned relatively unnatural values to “know”:

Why have a notion of knowledge that works in the way described?  
(Not a compulsory question. Enough to observe that we do have it.)  
But I venture the guess that it is one of the messy short-cuts — like satisficing, like having indeterminate degrees of belief — that we resort to because we are not smart enough to live up really high, perfectly Bayesian, standards of rationality. [...] If you doubt that the word ‘know’ bears any real load in science or metaphysics, I partly agree. The serious business of science has to do not with knowledge *per se*; but rather, with the elimination of possibilities through the ev-

idence of perception, memory, etc., and with the changes that one's belief system would (or might or should) undergo under the impact of such eliminations. Ascriptions of knowledge to yourself and others are a very sloppy way of conveying very incomplete information about the elimination of possibilities. It is as if you had said:

The possibilities eliminated, whatever else they may also include, at least include all the not- $p$  possibilities; or anyway, all of those except for some we are presumably prepared to ignore just at the moment.

(Lewis, 1996, 563)

Similarly, on the template mixed semantics above, saying that one knows  $p$  amounts to say that the possibilities one eliminates include those not- $p$  possibilities that are similar to their case *and* those that we take into consideration at the moment.

That the semantics ascribes an unnatural value to “know” is a reason to doubt the semantics itself. All things being equal, natural language words tend to have natural semantic values (Lewis, 1983, 1984; see Hawthorne, 2006, 205–7 and Sider, 2011, §3.2 for recent discussions). We are disposed to apply the word “cow” to cows and animals that look like cows. But “cow” only applies to cows because cows form a more natural group than cows and the animals that look like them.

Lewis gives two answers to those doubts. First, we may just observe that “know” happens to have such a semantics. For instance, we may grant that the template invariant semantics is more natural than the mixed contextualist one but argue that the latter is better supported by how competent speakers react to scep-

tical arguments. If the unnatural semantics is more than a accidental quirk of English, there must be some explanation for why those unnatural values are picked up. But it is not *a priori* excluded. Second, Lewis sketches just such an explanation:

The only excuse for giving information about what really matters in such a sloppy way is that at least it is easy and quick! But it *is* easy and quick; whereas giving full and precise information about which possibilities have been eliminated seems to be extremely difficult, as witness the futile search for a ‘pure observation language’. If I am right about how ascriptions of knowledge work, they are a handy but humble approximation. They may yet be indispensable in practice, in the same way other handy and humble approximations are. (Lewis, 1996, 563)

The idea is this: the natural properties that are close to the relatively unnatural semantic values we have are too difficult for us to track or less relevant to our interests. Hence naturalness is overridden by accessibility and usefulness. Similar constraints arise in other languages and can explain a convergence on similar meanings.

Lewis’s answers do not put the doubts to rest. He merely considers one competitor to his mixed semantics: a language that describes a subject’s evidence by stating the exact set of cases that are internally like it.<sup>18</sup> But there are others.

---

<sup>18</sup>See Lewis (1996, 553): “it is easier to list some of the propositions that are true in *all* the uneliminated, unignored possibilities than it is to find propositions that are true in *all and only* the uneliminated, unignored possibilities.”

One is the template invariant semantics above. Others are relational contextualist semantics — discussed in the next section. They do not appear less accessible or less useful than the mixed value for “know”. Consider the template invariant semantics, for instance. If we are able to track which not- $p$  possibilities are eliminated among the similar ones and those that we take seriously, presumably we are able to track which not- $p$  possibilities are eliminated among the similar ones. As far as usefulness is concerned, one may think that we typically care about which possibilities obtain *among those we are presently considering* and that the mixed value has an advantage on that score. But that is only true for knowledge ascriptions with interrogative complements. ‘ $S$  knows whether  $p$ ’ would tell us that  $S$ ’s evidence eliminates either the  $p$  or the not- $p$  possibilities we care about, without telling us which — so it is worth asking  $S$  herself. ‘ $S$  does not know whether  $p$ ’ would tell us that  $S$ ’s evidence eliminates neither — so it is not worth asking her. But on any semantics, ‘ $S$  knows that  $p$ ’ implies  $p$ . So ‘ $S$  knows that  $p$ ’ is no less efficient at ruling out not- $p$  possibilities speakers care about on an invariant semantics than it is on a mixed one.<sup>19</sup> By contrast, invariant relational semantics yield a knowledge predicate that is arguably more suited to evaluate what a subject ought to do. What matters to what a subject ought to do is the kind of situation she is in, not the kind of situation some speakers care about (see also [Hawthorne, 2004](#), 86–91). So usefulness does not clearly favour mixed values.

---

<sup>19</sup>Negation of knowledge ascriptions — ‘ $S$  does not know that  $p$ ’ — would have an added usefulness on the mixed semantics if they did not presuppose that  $p$ . But they do ordinarily presuppose that  $p$ , so on both semantics they already tell us that not- $p$  possibilities we care about are eliminated.

The present point is not to press these doubts, however. Let us temporarily grant that the mixed semantics is correct. Then what “know” refers to in most contexts is relatively unnatural. That suggests, as Lewis says, that what “know” refers to “bears no real load in science of metaphysics” (Lewis, 1996, 563, quoted above). It is barely worth theoretical study. By contrast, a natural property in the vicinity would be worth of theoretical study. But if there is such a property, it is the proper object of epistemology. And if it is close enough, epistemologists may call it “knowledge” too. To take an analogy: in ordinary language, the word “fruit” does not apply to cucumbers. Yet as it used by biologists, it applies to cucumbers. The property biologists refer to is a natural one, worth of systematic study. The property we ordinarily refer to is unnatural though more useful for cooking. The ambiguity in “fruit” leads to occasional confusions — such as misguided pedantic claims that “tomato is a fruit, not a vegetable” — but it is mostly harmless. One could endorse similar views about ordinary *versus* scientific uses of “weight”, “time” or “colour”.<sup>20</sup> Now if our ordinary word “know” has an unnatural semantic value with close natural competitors, then epistemologists use, or at least should use, “know” with a slightly different meaning. If they do, what contextualists say about the semantics of “know” in its ordinary use is irrelevant to its use in epistemology.

The mixed semantics substantiate that worry. On that semantics, *know* refers

---

<sup>20</sup>To be clear: we need not endorse such views here. Naturalness considerations may be so strong that even in their ordinary uses those words refer to the natural properties. The only claim here is that *even if* naturalness is overridden in ordinary use, it nevertheless tends to prevail in theoretical contexts.

to a property evidence has when it eliminates a disjunctive set of error possibilities: those that are either taken into consideration of similar enough to a subject's own situation. An epistemologist may rightly wonder why she should be interested in it. She may do so even if that property is the one that "know" ordinarily expresses. Consider the question whether somebody knows that they have hands, for instance. On the mixed semantics, it may be answered negatively for two very different kinds of reasons: because some error possibility is very similar to their situation or because we take seriously some error into consideration. Our epistemologist would naturally wonder which. She may introduce a term, "inv-knows", to say that one's evidence in a case eliminates error possibilities that are similar enough to that case. An obvious advantage of the term is that it is not context-sensitive: that spares her the worry that she equivocates when reasoning with it or evaluating the claims her colleagues make using it. Moreover, pretty much everything there is to say about "know" on the mixed semantics derives from (a) what people inv-know and (b) which possibilities speakers take into consideration in which contexts. Since (b) is mostly a matter of pragmatics and psychology, she may regard (a) as the core epistemological question. She may in particular wonder whether we inv-know anything. She may find that the most interesting sceptical arguments are those that equally apply to inv-knowledge (see for instance [Sosa, 2000](#), 6). In fact, it is tempting for our epistemologist to drop the prefix, call that property "knowledge" and stop worrying about what the ordinary folk say.

To sum up, mixed semantics are relatively unnatural. That itself raises doubts about their plausibility. But independently of these doubts, that raises a worry



for their relevance to epistemology. If the mixed semantics was correct, ordinary uses of “know” would refer to properties of little theoretical interest. There would be reason for epistemologists to use “know” to refer to theoretically interesting properties in the vicinity instead. Mixed contextualist semantics would then be mostly irrelevant to epistemology.

## 4 Relational semantics

### 4.1 Attention to respects of resemblance

We now turn to relational semantics. We first consider an account in the spirit of Lewis’s original account that relies on the idea that sceptical arguments exert their forces by drawing attention to *respects of resemblance* — instead of possibilities. The account is relational but raises the usual problems of Attention-based accounts. However, no alternative in terms of Taking Seriously is available here.

The proposal is this. We simply replace the Rules of Attention and Resemblance by the single following Rule:

**Rule of Attention to Respects of Resemblance** For any context  $C$ , if speakers in  $C$  attend to a respect of resemblance  $R$  between case, then any case  $c^*$  that resembles a case  $c$  in respect  $R$  is relevant to  $c$  in  $C$ .

The account draws on one understanding of Lewis’s appeal to “salient resemblance” (Ichikawa, 2011a, 395–6). The other Rules can be maintained in their relational versions. Alternatively, they can be re-framed as constraints on which

respects of resemblance are attended to. For instance, the Rule of Actuality may be recast as the claim that identity is a respect of resemblance that is always attended to.

The account has two advantages over Lewis's. First, it is somewhat more elegant. Lewis introduced an *ad hoc* exception to his Rule of Resemblance to prevent scepticism from prevailing in all contexts: resemblance in terms of the subject's evidence is not sufficient for relevance (Lewis, 1996, 556). But he had to rely on the Rule of Attention to allow scepticism to prevail in some contexts. The revised account dispenses with the exception and the separate rule. In ordinary contexts, speakers do not attend to internal resemblances but only to external or overall respects of resemblance. Sceptical scenarios exert their force by drawing speaker's attention to internal similarities.<sup>21</sup> Second, the account avoids asymmetry and irrelevance worries. In our asymmetry case, speakers attend to the possibility that the animal in the pen is a disguised mule. By doing so, they attend to the fact that actuality and that case resemble each other in terms of the way the animal looks. The case in which the wolves are fake resembles actuality in the same respect. Provided Alice's recognitional abilities are not better with wolves than with zebras, both (3) and (4) fail in the context. The relevance worry is avoided insofar as the respects of resemblance attended to are natural.

To secure these advantages, the account makes controversial assumptions about attention. It assumes that speakers do attend to respects of resemblances, that they

---

<sup>21</sup>External and overall respects of resemblance may need to be granted salience in *all* contexts in order to ensure that Gettier cases are never classified as knowledge cases. See Cohen (1998, 297–8) and Ichikawa (2011a, 396–7) for further discussion.

attend to fairly natural respects of resemblance, that in the zebra/wolf case they do attend to resemblance in terms of appearance of the animals. Each assumption can be challenged. Perhaps the most questionable assumption is that in ordinary contexts we do not attend to resemblances in terms of experience and memories. On the semantics, the truth of knowledge ascriptions depends on what goes on classes of cases that are perfectly similar in terms of experience and memories. It would thus be natural for speakers who make those ascriptions to attend to that respect of resemblance. An *ad hoc* exception may still be needed. If, in addition, Lewis's other Rules are recast as constraints on attention, even more controversial assumptions are needed and the resulting notion of attention may lose touch with the ordinary psychological one.

The main liability of the view, however, is its reliance on attention itself. Attending to sceptical scenarios does not always generate sceptical intuitions. Some speakers may watch a film like *The Matrix*, in which people are deceived in an brain in a vat fashion, and feel no pressure to deny knowledge to themselves and people around them. The revised account has a little more room than Lewis's. The film watchers need not attend to the internal similarities between their case and the deceived fictional characters' case, or between the latter and the normal cases in which these characters take themselves to be. They need not; but they may and yet fail to feel any sceptical pressure. As with Lewis's original account, a defender of the view could reply that the intuitions of these speakers are off track. But it is dialectically awkward for a contextualist to do so.

One cannot avoid the problem by simply replacing attention with some attitude

of taking something seriously. For respects of resemblance are not the kind of thing that can be taken seriously, as we saw in section 3.1.

## 4.2 Taking possibilities into consideration and comparative similarity

While we cannot adapt the attitude of taking seriously to respect of resemblances, we can go back to the Taking Possibilities Into Consideration diagnosis. The idea can be accommodated within relational semantics, but it is not trivial nor unproblematic to do so.

As we saw, attitudes of taking possibilities into consideration only yield rigid parameters — see 3.1. But there is a way to turn any rigid Rule in a relational one. We can use it to build a relational semantics on the basis of attitudes of taking seriously. The general idea is the following.<sup>22</sup> We assume a *comparative similarity* relation among cases: *a is at least as similar to b as c is to d*. The relation is necessary and sufficient to talk of degrees of similarity (Williamson, 1988): any given pair of cases  $\langle a, b \rangle$  can be mapped to a degree representing the similarity of *a* to *b*. The pair can act as a yardstick with which we pick up all pairs of cases that are similar to each other to that degree or more: the pairs  $\langle x, y \rangle$  such that *x* is at least as similar to *y* as *a* is to *b*. So given a pair of two cases  $\langle a, b \rangle$ , we can derive a *relation* between any two cases, namely the relation of one being at least as similar to the other as *x* is to *y*. By that method we can extract

---

<sup>22</sup>I am grateful to an anonymous referee for suggesting this.

a candidate relevance relation from various pairs of particular cases at play in a context, such as the case of the context itself, cases that are taken seriously, cases under discussion, or cases of evaluation.

Here a proposal along these lines. We use as a yardstick for our relevance relation the degree of similarity between the case of the context and cases taken seriously:

**Relational Rule of Taking Seriously (anchored in the speaker's case)** For any context  $C$  and case  $c$ : if a case  $c^*$  is as similar to  $c$  as some case taken seriously by the speakers of  $C$  is similar to their own case, then  $c^*$  is relevant to  $c$  in  $C$ .

The proposal implements the idea that as speakers take seriously that are very dissimilar to their own situation, their standards for “know” rise. The other Rules, construed relationally, can either be independently maintained or built into the notion of comparative similarity. In the latter option, the Rule of Actuality becomes the idea that any case is as similar to itself as any case is to any other. The first Rule of Reliability becomes the idea that any case in which a reliable process succeeds is *ceteris paribus* more similar to cases in which the process succeeds than to cases in which the process fails. And so on. Since we do not assume that the similarity relation in question is attended to, no worry of psychological plausibility arises. Under the mild assumption that in each context, at least one case is taken seriously, the Relational Rule of Taking Seriously would be the only one we need. The account is simple, it avoids the asymmetry and irrelevance worries and the pitfalls of attention.

### 4.3 The speakers in a vat problem

Though attractive, the account faces what we may call the *Speakers in a Vat problem*. Contextualists who appeal to the Taking-Seriously idea roughly think that speakers are in sceptical contexts when, and only when, they take sceptical scenarios seriously. Mixed semantics yield that result, albeit in an asymmetrical fashion. But the proposal above does not. For assume the proposal is correct and suppose that our speakers happen to *be* in a sceptical scenario. On the one hand, merely taking *ordinary* scenarios seriously is sufficient to unwittingly land them in a sceptical context. For their actual case is very dissimilar to the ordinary scenarios. On the other hand, taking *sceptical* scenarios like their own seriously will not be enough to raise their standards. For their case is similar to these scenarios. One may think that that is without consequence: since these speakers are in sceptical scenarios, they fail to satisfy “know” by any standard anyway. But such speakers may apply “know” to situations other than their own. *Prima facie* there is no reason why they could not truly utter, for instance: “If an ordinary person in a normal environment perceives an apple in front of them, they know that there is an apple in front of them”. Yet the proposal entails that insofar as they take ordinary possibilities seriously, they cannot. So the proposal fails to uphold the contextualist intuition that speakers in sceptical scenarios can use “know” in a non-demanding way. Relatedly, it must endorse a dialectically awkward error theory. For speakers in such scenarios turn out to be deceived, not only about the truth of their knowledge ascriptions, but also about their truth-conditions.

The Speakers in a Vat problem affects other proposals. A first anchors the

relational Rule of Taking Seriously to cases *under discussion*. In a given context, a case is relevant to a target case if it is as similar to that target case as *some case discussed in that context* is to some case taken seriously in that context. Again, we get the unwanted result that speakers who only take ordinary scenarios seriously but unwittingly discuss a sceptical situation (their own, for instance) are in sceptical contexts. Three other proposals fare better at first sight. The first is to anchor the Rule to cases *that the speakers believes to be under discussion*. The second is to anchor the the Rule to *cases that speakers believe might be actual*. The third is to take as yardstick the *spread of cases taken seriously*: a case is relevant to a target case if it is as similar to it as some two cases taken seriously are to each other. Insofar as our Speaker in Vat believes that only ordinary cases might be actual and only ordinary cases are under discussion, and insofar as they do not take sceptical scenarios seriously, they are in a non-sceptical context, as desired. But on closer inspection, our Speaker in a Vat *does* take a sceptical scenario seriously: namely, their own. Of course, they do not take it seriously under the mode of presentation “the case in which one is deceived in such-and-such a way”. But they do take it seriously under the mode of presentation “my case”. Since a sceptical case is taken seriously and very dissimilar from the anchor case (on the two first options) or other cases taken seriously (on the third one), the context is sceptical. To avoid the consequence, one may try to restrict the notion of “taking seriously” to cases specified in a descriptive manner. But it is doubtful that such a restriction is available or well-motivated. Even if such a restriction could be devised, the three proposals would run into trouble with admittedly *recherché* versions of

the Speaker in a Vat problem. A speaker who is convinced to be in a sceptical scenario, and only discusses and takes seriously such scenarios, should be in a sceptical context. But the three proposals fails to ensure she is.

Granted, intuitions about the Speaker in a Vat are less than compelling. Our aim is only to show that integrating the Taking Possibilities Seriously diagnosis of sceptical arguments into a relational semantics is far from trivial. First, the idea can be implemented in distinct ways. All share the idea that the standards for “know” are a function of the degree of similarity between some case taken seriously and some anchor case. But there are various options for anchor cases: the case of the context, the case under discussion, cases that speakers take to be candidates for actuality, other cases taken seriously, and so on. Lewisians have to decide which option is the best and why. Second, the Speaker in a Vat problem raises doubts as to whether the shared idea is true. The problem rests on contexts which are sceptical but where there is no great degree of dissimilarity between cases taken seriously and the candidate anchor cases. If there are such contexts, Lewisian relational semantics that integrate comparative similarity and the Taking Possibilities Seriously diagnosis fail.

#### **4.4 Relational semantics without Rules?**

Alternatively, Lewisians can endorse a relational semantics without specific Rules dictating how contexts affect the relevance relation. As we noted, [Ichikawa \(2011a, 387\)](#) thinks that Lewis’s analogy with quantifiers is on the right track but that no simple system of Rules “can be better than rather approximately correct”. He



does not explicitly state a relational semantics. But we can formulate one. We draw again on the idea of degrees of similarity:

**Similarity-Threshold relational semantics** For any context  $C$  and cases  $c, c^*$ :

$c^*$  is relevant to  $c$  in  $C$  iff the degree to which  $c^*$  is similar to  $c$  is above the threshold set by  $C$ .

More sophisticated proposals would not only allow the degree of similarity to vary with context, but also the dimension of similarity itself. We can leave such refinement aside for present purposes.

By relying on degrees of similarity, the view avoids the asymmetries and philosophical relevance issues. But can it provide an explanatory diagnosis of sceptical paradoxes?

The barest option is to simply state that the threshold is low in whatever context we deem to be sceptical. The view is not obviously wrong, but it is unsatisfactory. It does not say why sceptical arguments are more successful than, say, low spirits, in generating sceptical judgements.

At this juncture [Ichikawa \(2011a, 391\)](#) himself appeals to the Taking Seriously idea. Sceptical arguments draw attention to sceptical possibilities, and thus tend to make them relevant, which results in sceptical contexts. He merely refuses to turn the Taking Seriously idea into a hard-and-fast Rule (387). But here we face a dilemma.

Either we integrate the Taking Seriously idea with the degree of similarity semantics. That is in the spirit of the quantifier analogy. But we have seen that this runs into the Speakers in a Vat problem.

Or we simply state that the threshold of similarity lowers when sceptical scenarios are taken seriously. That delivers the desired result that contexts are sceptical when and only when sceptical scenarios are taken seriously. But the claim is *ad hoc* and unexplanatory. We are not told anything that explains why sceptical scenarios lower the similarity threshold. Rather, it is clear from the semantics that a low threshold of similarity makes for sceptical contexts, and it is independently assumed that taking sceptical scenarios seriously makes for sceptical contexts. The proposal simply postulates a brute connexion between the two but fails to explain it.

Summing up, several options are open for relational Lewisian semantics. Some are unpalatable, some need yet to be articulated. They may give up on explaining the alleged contextual effect of sceptical arguments. They may bite the bullet on the Speakers in a Vat problem. Alternatively, they may look for another way of integrating the Taking Possibilities Into Consideration diagnosis with relational semantics. Or they may look for another diagnosis of sceptical arguments.

## **5 Conclusion**

Lewis's core idea is that 'knows  $p$ ' is true of a case in a context just if one's evidence in that case eliminates relevant not- $p$  cases. When fleshing out the idea in a proper semantics, Lewisians face a choice. They can take context to select relevant cases rigidly, relationally, or in a mixed way. While the rigid option is a natural reading of Lewis's text, it is hopeless because it violates factivity.

Lewisians are left with mixed and relational semantics.

Their choice between the two depends in particular on whether they endorse the idea that sceptical arguments exert their force by drawing possibilities into consideration, and if they do, on how they intend to implement it. The most straightforward way to do so yields a mixed semantics. But on some natural assumptions, the resulting view predicts implausible asymmetries in the truth-value of some knowledge ascriptions. More generally, mixed semantics raise a worry that the epistemologist's use of "know" significantly differs from the ordinary one, thereby making the semantics mostly philosophically irrelevant.

Relational semantics are obtained in several ways. One draws on the idea that sceptical arguments exert their force by drawing attention to certain respects of resemblance. Because it relies on attention, the view is excessively liberal towards scepticism. Another is to combine the idea of taking possibilities into consideration with a comparative similarity relation. There are many ways to do so, but the most obvious ones run into trouble when considering speakers who are themselves in sceptical scenarios.

Lewisians may brush aside some of the problems or attempt to provide a new version of their view. But a fully satisfactory version of Lewisian contextualism has yet to be formulated.

## References

- Bach, K. (2005). The emperor new "knows". In Pryer, G. and Peter, G., editors, *Contextualism in Philosophy*, pages 51–90. Oxford University Press.
- Blome-Tillmann, M. (2007). Contextualism and the epistemological enterprise. *Proceedings of the Aristotelian Society*, 107:387–94.
- Blome-Tillmann, M. (2009). Knowledge and presuppositions. *Mind*, 118:241–294.
- Blome-Tillmann, M. (2012). Contextualism and the problem of known presuppositions. In Brown, J. and Gerken, M., editors, *Knowledge Ascriptions*, pages 104–. Oxford University Press.
- Cohen, S. (1998). Contextualist solutions to epistemological problems: Scepticism, gettier and the lottery. *Australasian Journal of Philosophy*, 76(2):289–306.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13:57–89.
- Douven, I. (2005). Lewis on fallible knowledge. *Australasian Journal of Philosophy*, 83(4):573–580.
- Feldman, R. (2001). Skeptical problems, contextualist solutions. *Philosophical Studies*, 103(1):61–85.

- Hawthorne, J. (2002). Lewis, the lottery and the preface. *Analysis*, 62:242–51.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford University Press.
- Hawthorne, J. (2006). Epistemicism and semantic plasticity. In *Metaphysical Essays*, pages 185–210. Oxford University Press.
- Hintikka, J. (1962). *Knowledge and Belief*. Cornell University Press.
- Ichikawa, J. (2011a). Quantifiers and epistemic contextualism. *Philosophical Studies*, 155:383–98.
- Ichikawa, J. (2011b). Quantifiers, knowledge and counterfactuals. *Philosophy and Phenomenological Research*, 82:287–312.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8:339–359.
- Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, 61(December):343–377.
- Lewis, D. (1984). Putnam's paradox. *Australasian Journal of Philosophy*, 62:221–36.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74:549–567.
- Oakley, I. T. (2001). A skeptic's reply to lewisian contextualism. *Canadian Journal of Philosophy*, 31:309–332.

- Pryor, J. (2001). Highlights of recent epistemology. *British Journal of Philosophy of Science*, 52(1):95–124.
- Schaffer, J. (2001). Knowledge, relevant alternatives and missed clues. *Analysis*, 61(271):202–8.
- Schaffer, J. (2004). Skepticism, contextualism, and discrimination. *Philosophy and Phenomenological Research*, 69(1):138–155.
- Schaffer, J. (2005). What shifts? thresholds, standards, or alternatives? In Pryor, G. and Peter, G., editors, *Contextualism in Philosophy*, pages 115–130. Oxford University Press.
- Sider, T. (2011). *Writing the Book of the World*. Oxford University Press.
- Sosa, E. (2000). Skepticism and contextualism. *Philosophical Issues*, 10:1–18.
- Stalnaker, R. (1984). *Inquiry*. MIT Press.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.
- Stanley, J. and Szabó, Z. G. (2000). On quantifier domain restriction. *Mind and Language*, 15(2&3):219–61.
- Vogel, J. (1999). The new relevant alternatives theory. *Noûs*, 33 (s13):155–80.
- Williams, M. (2001). Contextualism, externalism and epistemic standards. *Philosophical Studies*, 103:1–23.

Williamson, T. (1988). First-order logics for comparative similarity. *Notre Dame Journal of Philosophical Logic*, 29:457–81.