Index, context, and the content of knowledge[*]

Brian Rabern

`University of Edinburgh`

draft | 09/04/2016 | 12:37

The verb 'knows' is often treated as an intensional operator, and the logic of knowledge is often modelled with a modal operator *K*—a quantifier over ways that things could be compatible with a subject's knowledge (or evidence). 'Knows' is also often taken to be context-sensitive in an interesting way. What 'knows' means seems to be sensitive to the epistemic features of the context, e.g. the epistemic standard in play, the set of relevant alternatives, etc. There are standard model-theoretic semantic frameworks which deal with both intensional operators and context-sensitive expressions. The basic elements of these frameworks were developed by Montague (1968), Scott (1970), and Lewis (1970), and then received more sophisticated renditions and philosophical interpretations in Kaplan (1989a) and Lewis (1980). In this chapter, we provide a brief overview of the various moving parts of these frameworks, the roles of context and index, the need for double indexing, and the relationship between semantic value and content. With the best version of the standard framework explicated we then return at the end to the treatment of 'know' as a context-sensitive intensional operator (contrasting it with an invariantist treatment).

1   The semantics of parameter shifting and sensitivity

Let's begin with intensional model theory. Modal logic and tense logic saw great advances in mid-twentieth century works such as Carnap (1947), Prior (1956), Hintikka (1957), Kripke (1959), and Montague (1960).[1] The general upshot of this work was that the semantics of intensional constructions can be analysed (model-theoretically) in terms of quantification over parameters or "points" of some kind, e.g. possible worlds, times, etc.[2] Where □ is the modal necessity operator and F is the temporal futurity operator, the respective semantic clauses in terms of quantification over such "points"—worlds w and times t—are standardly given as follows.[3]

- $[\![\Box\phi]\!]^w = 1$ iff for all worlds w' (accessible from w), $[\![\phi]\!]^{w'} = 1$.

- $[\![\mathsf{F}\phi]\!]^t = 1$ iff there is a time $t' > t$ such that $[\![\phi]\!]^{t'} = 1$.

---

[*] For the *Routledge Handbook of Epistemic Contextualism*, Ichikawa, J. (ed.). Thanks to Adam Carter and Derek Ball for helpful comments.

[1] See Copeland (2002) for a detailed account of the genesis of possible world semantics.

[2] A further important feature of this general approach to intensional semantics is the inclusion of a binary *accessibility relation* on the space of points.

[3] Following the notational convention of Dana Scott, we let the equation $[\![\phi]\!]^i_{\mathfrak{A}} = 1$ mean "ϕ is true at point i relative to model $\mathfrak{A}$", but we suppress the model henceforth. See Scott (1970: 150-151) and Rabern (2016).

An interesting feature of these clauses is that a sentence is not just true or false relative to a model (as in e.g. propositional logic) but also relative to a *point of reference* (e.g. a world or a time) within a model.[4] Thus the semantics is "intensional" or non-extensional. This feature, however, isn't altogether novel to modal and tense logic, since the Tarskian semantics for quantification (given relative to a sequence of individuals g) is already "non-extensional" in this sense.[5]

- $[\![\forall \alpha \phi]\!]^g = 1$ iff for all sequences g' (that are α-variants of g), $[\![\phi]\!]^{g'} = 1$.

By the 1960's, theorists began applying the resources of this type of semantics to the study of languages involving "context-sensitivity" or "indexicality"—the phenomenon whereby the *meaning* of an expression depends on the context of use. Richard Montague (in Montague 1968 and Montague 1970a) called such languages "pragmatic languages" and suggested that a systematic treatment could be achieved by extending the tools of possible world semantics.

> It seemed to me desirable that pragmatics should at least initially follow the lead of semantics—or its modern version, model theory, which is primarily concerned with the notions of truth and satisfaction (in a model, or under an interpretation). Pragmatics, then, should employ similar notions, though here we should speak about truth and satisfaction with respect not only to an interpretation but also to a context of use. (Montague 1970a: 1)

With this approach in mind early theorists, e.g. Montague (1968), Scott (1970), and Lewis (1970), proposed that we simply expand the points of reference (or "indices") used for languages with modal and tense operators to include the relevant contextual coordinates.[6] For example, Scott advised as follows:

> For more general situations one must not think of the [point of reference] as anything as simple as instants of time or even possible worlds. In general we will have
>
> $$i = (w, t, p, a, \dots)$$
>
> where the index i has many coordinates: for example, w is a world, t is a time, $p = (x, y, z)$ is a (3-dimensional) position in the world, a is an agent, etc. All these coordinates can be varied, possibly independently, and thus affect the truth-values of statements which have indirect references to these coordinates. (Scott 1970: 151)

Consider a sentence that contains the first person pronoun.

(1)   I am a spiteful man

---

[4]Scott (1970) suggests, "One could call [them] *points of reference* because to determine the truth of an expression the point of reference must be established... Maybe *index* is just as good a term, though it seems to me to make them sound rather insignificant" (150).

[5]First-order logic just is a modal logic—its a multi-dimensional modal logic. See discussion in Blackburn et al. (2001).

[6]Davidson (1967) also suggested that for natural language semantics, truth should be relativized to times and persons in order to accommodate tense and demonstratives (see Davidson 1967: 319–320). Also notable in this regard is the "egocentric logic" developed in Prior (1968a): "If I say, not 'Brown is ill' but 'I am ill', the truth of this depends not only on when it is said but on who says it. It has been suggested, e.g. by Donald Davidson 1967 that just as the former dependence has not prevented the development of a systematic logic of tenses, so the latter should not prevent the development of a systematic logic of personal pronouns." (193).

Since the truth of this sentence depends crucially on who utters it, truth must be relativized to agents a in addition to worlds w and times t.

- $[\![(1)]\!]^{w,t,a} = 1$ iff a is in the extension of 'spiteful man' at world w and time t.

The essential idea was to generalize the techniques of intensional semantics to sentences containing context-sensitive expressions: Given that theorists already had the points of reference (i.e. worlds, times, and variable assignments) used for the semantics of modality, tense, and first-order quantification, a straightforward way to incorporate context-sensitivity was to expand the reference points to include various contextual parameters (e.g. speaker, place, addressee, demonstrata, etc.). A model-theory that made use of these expanded indices was thought to afford a formal unified treatment of both intensionality and indexicality. Sentential truth is *sensitive* to a parameter and certain sentential operators *shift* the parameter.

POINT OF REFERENCE: i

EXTENSION: $[\![.]\!]^i$
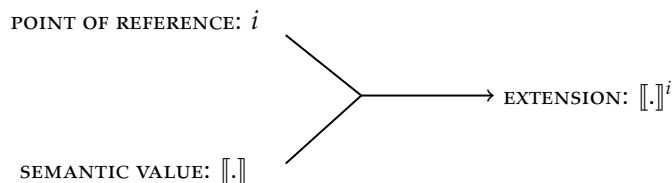
SEMANTIC VALUE: $[\![.]\!]$

Kaplan (1989) charged that the notion of a "point of reference" employed by early theorists blurred an important conceptual difference between "context of utterance" and "circumstance of evaluation".[7] He insisted on a two-step semantic procedure, which resolved all context-sensitivity before proceeding, and which distinguished between two kinds of meaning, the *character* and the *content* of an expression. In Kaplan's semantic theory these two aspects of meaning play different roles: the content is the information asserted by means of a particular utterance, whereas, the character of an expression encodes what any utterance of the expression would have as content.

CHARACTERS $[\![.]\!]$: Contexts → (Circumstances → Extensions)

CONTENTS $[\![.]\!]^c$: Circumstances → Extensions

The general picture is this: the domain of the character function is a set C. Each $c \in C$ is a tuple (or determines a tuple) of content-generating parameters—these tuples are called "contexts of utterance". Character functions map contexts of utterance to contents. The content of an expression is itself a function from a set V to extensions. Each $v \in V$ is also a tuple of parameters, often assumed to be possible worlds (or worlds paired with times, locations, agents, etc.)—these are called "circumstances of evaluation". The resulting Kaplanian picture is as follows:

---

[7]Montague, however, seems to have kept a division between the elements of the point of reference that were part of what he called the "context of use" and the element that was the "possible world" (see Montague 1970b, pp. 379-380; Cf. Israel and Perry 1996: 7-8). It is unclear what role, if any, this distinction actually played in Montague's semantic theory—but one might read into it a proto-character/content distinction, since Montague seems to identify intensions with Fregean senses.

CIRCUMSTANCE: v

CONTEXT: c $\qquad$ EXTENSION: $[\![.]\!]^{c,v}$

CONTENT: $[\![.]\!]^{c}$

CHARACTER: $[\![.]\!]$
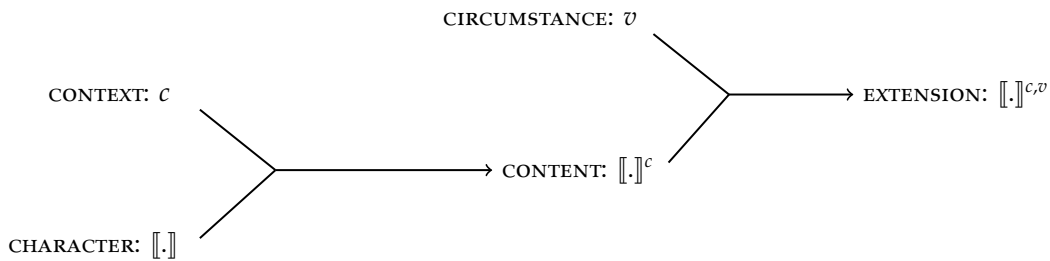
There are two independent, often conflated, reasons for Kaplan's insistence on this two-step procedure:[8]

- a linguistic motivation stemming from the compositional interaction of intensional operators and indexicals (i.e. double or multiple indexing)

- a pragmatic motivation stemming from the notion of assertoric content ("what is said") and its broader role in communication.

Let's consider each in turn.

1.1 Double indexing

The motivation from the compositional interaction of intensional operators and indexicals doesn't actually motivate the character/content distinction. Instead it simply calls for points of reference to be doubly indexed (cf. Lewis 1980). That a semantics for languages with indexicals embedded under intensional operators requires double indexing was first pointed out by Kamp (1971) with regard to tense logic.[9] The following sentences have the same truth-conditions:

(2)   It is raining.

(3)   It is raining now.

Originally this was taken to motivate a *redundancy* theory of 'now' (Prior 1968b). Letting N be the "now-operator" we could capture the redundancy with the following clause:

- $[\![\texttt{N}\phi]\!]^{t} = 1$ iff $[\![\phi]\!]^{t} = 1$

This accounts for the apparent "equivalence" between (2) and (3) but it also raises a difficulty, since (2) and (3) embed differently under other temporal operators:

(4)   It will be that case that it is raining.

(5)   It will be that case that it is raining now.

---

[8]See Stalnaker (forthcoming) for a nice discussion of this point. Cf. Rabern (2012).

[9]The ancestor to the 1971 paper on 'now' is Kamp (1967) "The treatment of 'now' as a 1-place sentential operator", which is eight pages of hand-written notes that Kamp presented to Montague's seminar on "pragmatics" at UCLA in 1967. These are stored in the Prior Archives at the Bodleian Library Oxford (Box 15). A copy is also available here: http://semanticsarchive.net/Archive/Tk3ZmEyN/

These have different truth-conditions, but given the semantics for N above we get the following equivalence (where F is defined as above) :

$$[\![\text{F rain}]\!]^t = [\![\text{F N rain}]\!]^t$$

Kamp showed that a singly-indexed semantics—a semantics that relativises to a point of reference with a single time parameter—cannot accommodate both (i) the data regarding the intuitive equivalence of between (2) and (3), and (ii) the difference in embedding behaviour between (2) and (3)—that is, the difference in semantic contribution that (2) and (3) make to more complex sentences such as (4) and (5).

> ... what we need is not just a definition of the notion: 'ϕ is true at i' but of the more complex notion: 'ϕ is true at i when occurring in a sentence uttered at j'. So our points of reference should be pairs $\langle i, j \rangle$ of moments of time, rather than single moments of time. (Kamp 1967)

A related way to see the need for double indexing is to consider the following sentence:

(6)   Everyone now alive will be dead.

The metalanguage truth-conditions can be represented as such:

$$[\exists t' : t' > t](\forall x(\texttt{alive}(x, t) \supset \texttt{dead}(x, t')))$$

And this requires two times in the point of reference, i.e. an assignment of times to *two* distinct temporal variables—for the two temporal variables in the embedded open formula.[10]

Such considerations—which stem purely from compositionality—motivate double (or multiple) indexing, but don't motivate Kaplan's character/content distinction and the two-step semantics.

1.2   Semantics and "what is said"

The other, more fundamental reason, for Kaplan's insistence on the two-step procedure concerns the relationship between semantics and the contents of assertion. Essential to Kaplan's two-step picture is a particular view about the division of theoretical labor between the components of a point of reference. There's a context, and there's a circumstance. Points of reference are treated as context-circumstance pairs. We need both, according to Kaplan, because we need our semantic theories to be able to capture the two different ways in which, when somebody says (for example), "You are a fool" or "I am a spiteful man", the truth of their utterance depends on the situation in which they say it—roughly, that the situation in which they say it influences both what was said, and whether whatever was said is true. Contexts play a content-generating role—resolving context-dependence in order to determine what's said—and circumstances play a content-evaluating role—they're the things of which what's said is either true or false.

---

[10]Vlach (1973) upped the ante by focusing on sentences such as the "past tense" version of (6): "Once everyone then alive would be dead". We can represent the logical form of this as follows: $[\exists t' : t' < t][\exists t'' : t'' > t'](\forall x(\texttt{alive}(x, t') \supset \texttt{dead}(x, t'')))$. Given the three temporal variables in the syntax this requires *three* times in the point of reference. With increasingly complex sentences involving further temporal embedding there is a need for further temporal parameters. See Cresswell (1990) for a detailed discussion.

point of reference = ⟨content-generators, content-evaluators⟩

In its content-generating role, the context provides all the various contextual parameters for the resolution of indexicals and other context-sensitive expressions—different people say different things, depending on who's speaking and who's being spoken to, with "You are a fool" and "I am a spiteful man". In its content-evaluating role, the circumstance provides various parameters appealed to in the semantics of intensional constructions. For example, in the evaluation of an utterance of "Necessarily, I am a spiteful man" the modal operator "Necessarily" checks whether, in every circumstance, things are as that particular utterance of "I am a spiteful man" represents things as being. If there is a circumstance of evaluation differing with respect to whether things are as the utterance, in context, represents things as being, then the utterance of the modalized sentence is false (and it's true otherwise).

According to Kaplan the two-step procedure is crucial, since a central task of a semantic theory is to tell us what sentences *say* in various contexts—what propositions or pieces of information do they express in a given context.[11] But contents in Kaplan's framework play an addition role, beyond just being the output of assertion. For Kaplan the content of a sentence also plays a compositional role of being the object of various sentential operators. Thus, contents are constrained depending on the operators of the language, to be the type of semantic entities that enter into compositional relations with those operators. For these reasons, Kaplan is led to endorse temporalism about propositions—the view that propositions can vary in truth value across times.

> If we built the time of evaluation into the contents. . . , it would make no sense to have temporal operators. To put the point another way, if what is said is thought of as incorporating reference to a specific time. . . it is otiose to ask whether what is said would have been true at another time. . . (Kaplan 1989a: 503)

Lewis insists that the two-step procedure isn't theoretically motivated—he contends that an equally good option is just to evaluate at both a context and index in one-step. Lewis emphasises that a theory of the first sort can be easily converted into one of the second and vice versa simply by currying or un-currying the functions.

The disagreement on this point between Kaplan and Lewis stems from their differing views on the role of assertoric content in the semantic theory. For Kaplan, content has a privileged role in the semantic theory proper: it is both "what is said" and the level of semantic value over which the composition rules can be defined. For Lewis, assertoric content is *post-semantic*. He agrees that "we can assign propositional content to sentences in context" and that "propositions have an independent interest as suitable objects for attitudes such as belief, and [illocutionary acts]" (37), but he doesn't build into the semantics proper an identification between assertoric content and sets of indices (i.e. semantic values in a context). That is, context and index are playing different roles in the two frameworks, for Lewis, in contrast to Kaplan, the "context" is not a sequence of content-generating parameters, and sets of indices are not propositions.

Lewis doesn't equate sets of indices with propositional content, since he doubts that one type of semantic entity can play both roles. In particular, he worries that the parameters that will be required in the indices to provide an adequate compositional semantics might result in sets of indices that are unfit to play the content role. Yalcin provides a nice summary of Lewis' worry as follows:

---

[11] "The idea of Content—the what-is-said on a particular occasion—is central to my account." (Kaplan 1989: 568)

> It is possible that, owing to the operators the language in question contains, the semantic value of a sentence relative to context must be some complicated intension, variable with respect to an array of parameters—say, parameters for world, time, location, standard of taste, orientation, standard of precision, state of information, etc. The details here will be a contingent matter concerning the particular architecture of the language in question. It has to do with what expressions (if any) are best semantically modeled as intensional operators. (Yalcin 2014: 23)

It was precisely for reasons stemming from "the particular architecture of the language" that Kaplan was led to the conclusion that content is variable with respect to a time parameter—this was due to temporal operators.[12] But since there are potentially further operators that would require further parameters, it seems unlikely that the resulting sets of circumstances will be apt as objects of assertion and the attitudes.

Kaplan's treatment of quantifiers provides a nice example of Lewis's point here (cf. Rabern 2013). The assignment function plays an essential role in the semantics—Tarskian semantics—for the quantifiers. Quantifiers attach to formulas and shift the assignment of values to variables. But if the content of the whole is determined by the content of the parts, then the assignment function is a parameter of the circumstance. Of course, Kaplan doesn't opt for this option, since then contents would be functions from worlds, times, and assignment functions, to truth values.[13] Contents, so conceived, seem ill-suited to play the content role. After all, contents represent what is said, the objects of attitudes such as belief and knowledge, and so forth. What sense does it make to say that what is said, or what is known, is true or false depending on an assignment function? This is exemplifies Lewis' point: one type of semantic entity can't play both the compositional and content roles.

Inherent in Lewis' discussion is a dilemma for any view that tries to identify the propositional content of a sentence in a context with the semantic value of a sentence a context.

1. *First horn*: Let the propositional contents of sentences (in context) be whatever they are according to the preferred account of the representational properties of mental states (belief-desire psychology) and the best account of assertion and communication. Whatever these entities turn out to be, it is very unlikely, given the particular architecture of the language, that they can also serve as the compositional semantic values of sentences.

2. *Second horn*: Let the semantic values of sentences (in context) be whatever they must be according to our best compositional semantic theory given the particular architecture of the language. Whatever these entities turn out to be, it is very unlikely that they will have independent interest in connection with the representational properties of mental states, and as the objects of assertion, belief and knowledge.

---

[12]Kaplan does notice that there is a tension between the functional notion of content in terms of compositionality and the traditional notion in terms of assertoric content, he says: "This functional notion of the content of a sentence in a context may not, because of the neutrality of content with respect to time and place, say, exactly correspond to the classical conception of a proposition. But the classical conception can be introduced by adding the demonstratives 'now' and 'here' to the sentence and taking the content of the result. I will continue to refer to the content of a sentence as a proposition, ignoring the classical use" (Kaplan 1989a, p. 504). He does not, however, seem to appreciate the implications of the gap between assertoric content and compositional value.

[13]For Kaplan the assignment function is a parameter of context (cf. Kaplan 1989b: 591). He defines the content of ϕ in the context c under the assignment f, which he symbolises with $\{\phi\}_{c,f}$, as follows (546) : $\{\phi\}_{c,f} = \{(w,t) \mid [\![\phi]\!]^{c,f,t,w} = 1\}$. The problem with this option is that it is inconsistent with Kaplan's ban on monsters—quantifiers are monsters (Rabern 2013).

Of course, someone can wiggle out of the dilemma by either insisting on a (non-standard) syntax/semantics according to which the compositional values are plausibly of independent interest as the objects of assertion or insisting on a theory of assertion and communication such that the entities appealed to there just happen to be the entities apt for compositionality. But Lewis' point is that we have no theoretical reason to expect such correspondence from the outset. Lewis sums up the situation as follows.

> It would be a convenience, nothing more, if we could take the *propositional content* of a sentence in a context as its *semantic value*. But we cannot. The propositional contents of sentences do not obey the composition principle, therefore they are not semantic values. (Lewis 1980: 39)

If there is no *a priori* constraint on semantic theorizing that a single type of entity plays both of these roles, we should not be worried when the demands of compositional semantics shape "meaning" in a way that is different from our preferred theory of attitude contents. This basic picture of the relationship between compositional semantics and the contents of attitudes has recently been developed and advocated in Yalcin (2007), Ninan (2010), and Rabern (2012). Stanley also endorses such a distinction in terms of Dummett's 'ingredient sense' and 'assertoric content' (Stanley 1997b and Stanley 2002).[14]

Nevertheless, one might insist that the things we say and the meanings of our words stand in a intimate and theoretically important relationship. After all, we utter words with certain meanings (and certain syntax) in order to say the things we say. Yet, this platitude does not call for the *identification* of the two notions—all it would call for is that the propositional content of a sentence in a context should be systematically *determined* by its semantic value. Lewis states:

> It is enough that the semantic value of a sentence in context should somehow determine the assignment of propositional content. And it does... we have the relation: sentence *s* is true at context *c* at index *i*. From that we can define the propositional content of a sentence *s* in context *c* as that proposition that is true at world *w* iff *s* is true at *c* at the index i_c^w that results if we take the index i_c of the context *c* and shift its world coordinate to *w*. (Lewis 1980: 37-38)

Just as there is a post-semantic definition of *truth-in-a-context* for a sentence given in terms of the sentence's semantic value, there is also a post-semantic definition of the *assertoric-content-in-a-context* of a sentence given in terms of the sentence's semantic value. To demonstrate this point assume that in addition to the context *c*, extension is relativised to an index consisting of an assignment, a time, and a world—thus points of reference are quadruples $\langle c, g, t, w \rangle$. The compositional semantics will recursively define $[\![.]\!]^{c,g,t,w}$. But we follow Lewis and don't build into the semantics proper an identification between assertoric content and sets of indices.[15] Instead we provide a postsemantic definition of content in terms of semantic values. For example, assuming a classic view of propositions, so that that contents are simply sets of worlds, we can define the content of a sentence ψ in a context *c*, which we write as $|\psi|^c$, as follows:

Assertoric content. $|\psi|^c = \left\{ w : [\![\psi]\!]^{c,g_c,t_c,w} = 1 \right\}$

---

[14] For background to Dummett's ingredient sense/assertoric content see Dummett (1973), Evans (1979), and Davies and Humberstone (1980). For more recent discussion of Lewis' semantic value/content distinction see Yalcin (2014).

[15] Arguably, linguists tend to follow Lewis in this regard. See Heim and Kratzer (1998) for some basic background, and Von Fintel and Heim (2011) for an extension of the framework to treatments of intensionality.

2 'Knows': index-shifting and context-sensitivity

The verb 'knows' has been given various semantic analyses in terms of the parameter sensitive semantic frameworks outlined in the previous sections. In natural language semantics it has been treated as an intensional operator in the manner of Hintikka (1962), and more generally the structure of knowledge has been explored in epistemic modal logic, where knowledge is represented by the modal operator K (see Holliday 2016 for an overview). 'Know' has also been given various analyses where it is sensitive to some contextual parameter, e.g. an epistemic standard, the relevant alternatives, a contrast classes, the question under discussion, etc. (Cohen 1986, DeRose 1995, Lewis 1996, Schaffer 2004). In what follows, we provide an analysis of 'knows' as an intensional operator, and then look at the differences between strict-invariantist, contextualist, and sensitive-invariantist treatments of such an operator.

2.1 'Knows' as an index operator

The idea that attributions of propositional attitudes such as belief, knowledge, memory, desire, etc. can be helpfully analysed as intensional operators goes back to Hintikka:

> . . . an attribution of any propositional attitude to the person in question involves a division of all the possible worlds. . . into two classes: into those possible worlds which are in accordance with the attitude in question and into those which are incompatible with it. The meaning of the division in the case of such attitudes as knowledge, belief, memory, perception, hope, wish, striving, desire, etc. is clear enough. For instance, if what we are speaking of are (say) a's memories, then these possible worlds are all the possible worlds compatible with everything he remembers. (Hintikka 1962: 91)

Hintikka's idea is to treat the semantics of propositional attitudes using off-the-shelf possible world semantics but vary the accessibility relations depending on the attitude. An operator like 'believes' is analysed as follows:

- $[\![\alpha \text{ believes } \phi]\!]^{c,w} = 1$ iff for all $w' \in W$ compatible with α's beliefs in w, $[\![\phi]\!]^{c,w'} = 1$

We could also write it in a fashion where it more explicitly appeals to an accessibility relation B^α by letting $B^\alpha(w', w)$ iff w' is compatible with α's beliefs in w. Likewise, 'knows' can be analysed as an intensional operator that quantifies over all the worlds compatible with the agent's *knowledge* (which again we could abbreviate as $K^\alpha(w', w)$):

- $[\![\alpha \text{ knows } \phi]\!]^{c,w} = 1$ iff for all $w' \in W$ compatible with α's knows in w, $[\![\phi]\!]^{c,w'} = 1$

Such a clause might not seem very informative—not, at least, to someone interested in the nature of knowledge. But it is not completely uninformative. One could investigate the logical properties of the accessibility relation induced by 'knows'. For example, K^α will be reflexive, whereas B^α will not be. One could also unpack the right-hand-side in a more informative way in order to reveal the connections to belief, truth, and justification, e.g., as follows:

- $[\![\alpha \text{ knows } \phi]\!]^{c,w} = 1$ iff

  (i) *Belief*: for all $w' \in W$ compatible with α's beliefs in w, $[\![\phi]\!]^{c,w'} = 1$, and

(ii) *Truth*: $[\![\phi]\!]^{c,w} = 1$, and

(iii) *Justified*: for all $w' \in W$ compatible with α's evidence in w, $[\![\phi]\!]^{c,w'} = 1$.

Of course, this particular semantics seems subject to Gettier counterexamples (Gettier 1963). This is just to demonstrate a way in which 'knows' could be analysed as an intensional operator—Gettier examples are not a special problem for treating 'knows' in such a fashion. For the remaining, we will not worry about adding extra "anti-Gettier" conditions. We will also set aside explicit mention of the belief and factivity conditions[16], and just focus on the core epistemic feature (with the understanding that the right-hand-side really has extra but suppressed sufficient conditions): α knows that ϕ iff α's evidence eliminates every possibility in which not-ϕ and [...]. Here we are only interested in certain abstract semantic features concerning the relationship between the meaning of 'knows' and the context of utterance.

2.2 'Knows' as context-sensitive

According to the analyses above 'knows' is an intensional operator, and it is *not* context sensitive. Consider the simplified clause (let A be a set of worlds, perhaps the set of all worlds).[17]

$[\![\alpha \text{ knows } \phi]\!]^{c,w} = 1$ iff for all $w' \in A$ compatible with α's evidence in w, $[\![\phi]\!]^{c,w'} = 1$

Since "contextualism" is a view about the assertoric content of a sentence in a context, namely the view that the content can vary with context, one can't read "contextualism" off of the semantic clause alone—we also need to appeal to the post-semantic definition of how content is determined by semantic value in a context.[18] Assume a version of the definition of assertoric content from above where the proposition expressed by an utterance of ψ in context c is just the set of worlds in which it is true given c:

$|\psi|^c = \left\{ w : [\![\psi]\!]^{c,w} = 1 \right\}$

Contextualism concerns the proposition expressed: does 'α knows ϕ' express the same assertoric content in every context or does the content vary across contexts. The simple clause for 'knows' is clearly not context sensitive, since for any two contexts c and c', we have (under the assumption that α and ϕ aren't context sensitive)

$|\alpha \text{ knows } \phi|^c = |\alpha \text{ knows } \phi|^{c'}$

Thus we have a non-contextualist treatment of 'knows'. Call this view "strict invariantism". On this view an utterance of the sentence "α knows ϕ" always expresses the same proposition—in

---

[16]Note that I'm not assuming that knowledge can be analysed into constituent components such as justification, truth, and belief—the discussion is intended to be compatible Williamson's 'knoweldge-first' project (Williamson 2002).

[17]We define 'knows' syncategorimatically for ease of presentation, and to avoid various issues with the internal syntax and compositional semantics of knowledge attributions.

[18]Not that there is complete convergence and agreement on the terminology, but I think the way I've set things up captures what is really at issue in the core debates. An extreme case worth considering is this: assume one insisted that the assertoric content of a sentence was always the diagonal, then content would never vary with context, no matter what semantics one proposed for 'knows'—is this compatible with contextualism?

every context it attributes *knowledge* to the agent.[19]

Contextualist views, on the contrary, hold that sentences containing 'knows' can express different propositions in different contexts. What evidence might support this claim? It's the same type of evidence that supports the claim that a sentence such as "I am hungry" (or "Hazel is tall") express different propositions in different contexts: linguistic evidence. Its is evidence concerning what we would or wouldn't say, or what would or wouldn't be true, in various situations. (See §1 for discussion of Data and Motivations.) In this case, if there are two contexts where utterances of the same sentence seem to have different truth conditions that is some evidence that the knowledge attributions are context sensitive. For example, consider this situation:

> *Situation 1.* Sam parked his car out front and came into the party. He had a few drinks and mingled with the guests. After a couple of hours we are gossiping with Sam about the party goers. The conversation turns to the increased number of micro-breweries in the city, and the story of one of our friends who just opened a new brewery in this neighbourhood. Sam suddenly grabs his jacket and heads toward the back door. We interrupt him to remind him that he came in through the front, not the back. Sam replies that he is not leaving he just going outside to smoke, and then he says
>
> (7) I know that my car is parked out front.

Contrast that with this case:

> *Situation 2.* Sam parked his car out front and came into the party. He had a few drinks and mingled with the guests. After a couple of hours we are gossiping with Sam about the party goers. The conversation turns to the increased rate of car thefts in the city, and the story of one of our friends who had a car stolen last weekend in this neighbourhood. Sam suddenly grabs his jacket and heads toward the back door. We interrupt him to remind him that he came in through the front, not the back. Sam replies that he is not leaving he just going outside to smoke, and then he says
>
> (8) I know that my car is parked out front.

Utterance (7) seems unproblematically true. But utterance (8) seems problematic—how can Sam know that his car is parked out front, if he doesn't know that his car hasn't been stolen? It seems that in situation 2, since the prospects of car theft have been made so salient the standards for Sam to count as having knowledge have been raised. There seem to be different epistemic standards operative in the two contexts, in such a way that the proposition expressed varies with the epistemic standard operative in the context.

A very basic kind of contextualism would have it that 'α knows ϕ' is true in c iff α meets the epistemic standard operative in c. An "epistemic standard" is just a place holder for how good one's epistemic position must be to count as having knowledge—we can construe it in quantitative

---

[19]Note that this is not to say that 'knows' is *completely* context invariant. A knowledge attribution might express different propositions at different contexts due to various other (non-epistemic) factors that we are currently ignoring. Stanley (2005) agrees, "As is standard, I assume that knowledge attributions, like other sentences, express propositions about particular times. One and the same knowledge attribution may express different propositions at different times, because the temporal element in the knowledge attribution is assigned different times in different contexts of use. So, knowledge attributions are context-sensitive, on my view, in the sense that different knowledge attributions are about different times. But this is a relatively innocuous sense of contextualism; it is true of all verbs that they are associated with temporal elements that receive different values at different times. On the semantic clause I have just given, there is no specifically epistemological sense in which knowledge attributions are context-sensitive."

terms such as the following: the level of justification required, or the amount of evidence required. Contextualism doesn't require that *evidence* plays a central role in the account of knowledge, it is compatible with more externalist epistemologies which might cash out the epistemic standard in terms of (e.g.) reliability. The exact nature of the epistemic standard is controversial and I will remain neutral here. I will simply assume that raising the epistemic standard requires ruling out more possibilities. For example, raising the amount of evidence required for knowledge can be understood as expanding the possibilities that must be ruled out by the evidence. So I will couch the discussion in a way that most resembles the brand of epistemic contextualism found in Lewis (1996) (see Schaffer 2015 for helpful discussion).

Lewis' gloss on the contextualist semantics was this:

> 'α knows ϕ' is true iff α's evidence eliminates every possibility in which not-ϕ—Psst!—except for those possibilities that we are properly ignoring.

That is, "α knows ϕ" is true in context c just in case α's evidence eliminates every not-ϕ possibility relevant in c. Since in different contexts different possibilities are relevant, different propositions are expressed. By contrast, on the invariantist view above the set of relevant possibilities A remained constant across contexts. Contrast the invariantist semantics with the contextualist semantics.

STRICT INVARIANTISM:

$[\![\alpha \text{ knows } \phi]\!]^{c,w} = 1$ iff for all $w' \in A$ compatible with α's evidence in w, $[\![\phi]\!]^{c,w'} = 1$

CONTEXTUALISM:

$[\![\alpha \text{ knows } \phi]\!]^{c,w} = 1$ iff for all $w' \in A_c$ compatible with α's evidence in w, $[\![\phi]\!]^{c,w'} = 1$

It is worth pausing here to see the key difference between the contextualist and the invariantist semantics. As I've construed things the difference comes out in a way that might look very insignificant. It comes down to whether or not the set of relevant alternatives A is determined by the context or not. For the invariantist the set of relevant alternatives is invariant—it is always just A.[20] Whereas for the contextualist the relevant set of possibilities A_c varies with c. The seemingly insignificant little c subscripted on the A can make a very significant difference: "A_c" means the set of worlds determined by the context c. Thus, "knows" requires ruling out all the possibilities in A_c, where this set varies with c—the worlds outside of A_c are "properly ignored". Notice that on the invariantist semantics there is no little subscripted c on A, thus A remains the same no matter the context.

In situation 1, the relevant set of alternatives A_{c_1} doesn't include worlds where the rate of car theft in the neighbourhood is significantly high, whereas in the context of situation 2—where the prospects of car theft have been made so salient—the relevant set of alternatives A_{c_2} has expanded to include worlds where the rate of car theft is significantly high. According to contextualism the variability data concerning knowledge ascriptions can be accounted for by this sensitivity to the contextually variable epistemic standards.

---

[20]One might assume that A is identical to the set of all worlds W, but this is not an assumption that the invariantist *per se* is committed to. In particular, an invariantist might insist that there are certain "hinge propositions"—perhaps there are possibilities outside of A but it is never required that our evidence rule out these possibilities, instead we have a special *entitlement* for these (see, e.g., Wright 2004).

There is also a brand of invariantism that purports to handle the variability data by appealing to other varying factors that the knowledge attributions are sensitive to.[21] Proponents of this view insist that it isn't the context of the knowledge attributions that matter—the relevant set of worlds doesn't depend on the context—instead what might vary is various factors concerning the agent of the attribution.[22] In this way the relevant set of words is sensitive to the alleged knower's situation—these factors can include things like her presuppositions or what error-possibilities she is considering and further factors that are often described as "non-epistemic" features such as "stakes" or "practical interests" (see e.g. Hawthorne 2004, Stanley 2005, and Chapter 19).

SENSITIVE INVARIANTISM:

$$[\![\alpha \text{ knows } \phi]\!]^{c,w} = 1 \text{ iff for all } w' \in A^\alpha_w \text{ compatible with } \alpha\text{'s evidence in } w, [\![\phi]\!]^{c,w'} = 1$$

For situation 1 the set of relevant alternative will be the set $A^\alpha_{w_1}$ which depends on Sam's particular situation in w_1, whereas for situation 2, the set of relevant alternative will be the distinct set $A^\alpha_{w_2}$ which depends on Sam's situation in w_2 (where he has just been considering certain theft possibilities, etc.). In this way, the sensitive invariantist claims to account for the variability data in a way that is similar to the contextualist, but while retaining invariantism.

In the case of first person knowledge ascriptions contextualism and sensitive invariantism are difficult to tease apart. But the views will make very different predictions about third person knowledge reports. For example, if we are watching Sam (who is oblivious to the recent increased theft rate) from afar it seems that the truth of our utterance of (9) can vary depending on what is salient in our context of attribution, e.g. the truth of our utterance of (9) can vary depending on whether or not *we* are discussing the recent increased theft rate.

(9)   Sam knows that his car is parked out front.

This has been taken to support contextualism over sensitive invariantism. This is not the place to assess contextualism versus sensitive invariantism, but these are the types of cases where the views make different predictions.

Another place where contextualism and subject sensitive invariantism might come apart is with respect to embeddings under various intensional operators, since for the invariantist, but not for the contextualist, shifts of the world can shift the set of relevant alternatives (cf. Stanley 2005: 106-116).

---

[21]There another important option—the "relativist" option—but I will suppress it in the main text in order to simplify the presentation. The relativist (see MacFarlane 2014: 187-190) will add to the index a parameter A for a set of relevant possibilities, and provide the following clause: $[\![\alpha \text{ knows } \phi]\!]^{c,w,A} = 1$ iff for all $w' \in A$ compatible with α's evidence in w, $[\![\phi]\!]^{c,w',A} = 1$. Without a view about how one moves from the compositional semantics to assertoric content, this clause is compatible with strict invariantism ($|\psi|^c = \{w : [\![\psi]\!]^{c,w,A} = 1\}$), contextualism ($|\psi|^c = \{w : [\![\psi]\!]^{c,w,A_c} = 1\}$), and sensitive invariantism ($|\psi|^c = \{w : [\![\psi]\!]^{c,w,A_w} = 1\}$). The real difference comes with the postsemantic definitions of truth and content. Macfarlane takes the "contents of attitudes to be 'relevant-alternatives-neutral'—that is, to have truth values that vary not just with the state of the world but with the relevant alternatives" (189). Thus, we get this relativist definition of content: $|\psi|^c = \{(w, A) : [\![\psi]\!]^{c,w,A} = 1\}$. MacFarlane explains as follows: "On this view, the relation 'knows' expresses does not vary with the context—there is just a single knowing relation—but the extension of that relation varies across relevant alternatives. As a result, it makes sense to ask about the extension of 'knows' only relative to both a context of use (which fixes the world and time) and a context of assessment (which fixes the relevant alternatives)" (189). The relativist definition of truth is relativized to both an utterance context c and an assessment content c': ψ is true in c and c' iff $[\![\psi]\!]^{c,w_c,A_{c'}} = 1$.

[22]Sometimes the distinction between sensitive-invariantism and contextualism is made by using the misleading terminology "subject contextualism" versus "attributor contextualism", but only the latter is a genuine form of *contextualism* as we have defined things here.

For example, take some attribution "α knows ϕ" that is true according to sensitive invariantism, since the stakes for α in w are so low. Embedding into the consequent of a counterfactual whose antecedent shifts to worlds with higher stakes, however, is predicted to vary the truth value of "α knows ϕ". In other words, a counterfactual such as the following is predicted (counterintuitively) to be true.[23]

(10)   If the stakes had been higher, then α wouldn't have known ϕ.

In such cases, and further cases involving temporal embeddings, our judgments don't seem to vary in the way the sensitive invariantist predicts. If so, that would be a reason to have the relevant alternatives tied to the context of utterance, instead of the world of the index (see Stanley 2005 Blome-Tillman 2009, and MacFarlane (2014: Chapter 8) for detailed discussion).

References

Blackburn, P., de Rijke, M. and Venema, Y.: 2001, *Modal Logic*, Cambridge: Cambridge University Press.

Blome-Tillman, M.: 2009, Contextualism, subject-sensitive invariantism, and the interaction of 'knowledge'-ascriptions with modal and temporal operators, *Philosophy and Phenomenological Research* **79**(2), 315–331.

Carnap, R.: 1947, *Meaning and necessity: A study in semantics and modal logic*, University of Chicago Press.

Cohen, S.: 1986, Knowledge and context, *Journal of Philosophy* **83**(10), 574–583.

Copeland, B.: 2002, The genesis of possible worlds semantics, *Journal of Philosophical Logic* **31**(2), 99–137.

Cresswell, M.: 1990, *Entities and Indices*, Kluwer Academic.

Davidson, D.: 1967, Truth and meaning, *Synthese* **17**(1), 304–323.

Davies, M. and Humberstone, L.: 1980, Two notions of necessity, *Philosophical Studies* **38**(1), 1–31.

DeRose, K.: 1995, Solving the skeptical problem, *Philosophical Review* **104**(1), 1–52.

Dummett, M.: 1973, *Frege: Philosophy of Language*, London: Gerald Duckworth.

Evans, G.: 1979, Reference and Contingency, *The Monist* **62**, 161–189.

Gettier, E. L.: 1963, Is justified true belief knowledge?, *Analysis* **23**(6), 121–123.

Hawthorne, J.: 2004, *Knowledge and Lotteries*, Oxford University Press.

Heim, I. and Kratzer, A.: 1998, *Semantics in Generative Grammar*, Blackwell Publishers.

Hintikka, J.: 1957, Modality as referential multiplicity, *Ajatus* **20**, 49–64.

---

[23]Imagine Hannah in a low stakes situation (in the scenario of Stanley 2005) saying "I know the bank is open, but if I had a bill coming due, then I wouldn't know".

Hintikka, J.: 1962, *Knowledge and Belief: An Introduction to the Logic of the Two Notions*, Cornell University Press.

Holliday, W. H.: 2016, Epistemic logic and epistemology, *in* S. Hansson and V. Hendricks (eds), *The Handbook of Formal Philosophy*, Springer.

Israel, D. and Perry, J.: 1996, Where monsters dwell, *in* J. Seligman and D. Westerståhl (eds), *Logic, Language and Computation*, Vol. 1, CSLI Publications, Stanford.

Kamp, H.: 1967, The treatment of 'now' as a 1-place sentential operator, *multilith*, University of California in Los Angeles.

Kamp, H.: 1971, Formal properties of 'now', *Theoria* **37**, 227–274.

Kaplan, D.: 1989a, Demonstratives, *in* J. Almog, J. Perry and H. Wettstein (eds), *Themes from Kaplan*, Oxford University Press, pp. 481–563.

Kaplan, D.: 1989b, Afterthoughts, *in* J. Almog, J. Perry and H. Wettstein (eds), *Themes from Kaplan*, Oxford University Press, pp. 565–614.

Kripke, S.: 1959, A completeness theorem in modal logic, *The Journal of Symbolic Logic* **24**(1), 1–14.

Lewis, D.: 1970, General semantics, *Synthese* **22**(1), 18–67.

Lewis, D.: 1980, Index, context and content, *in* S. Kanger and S. Ohman (eds), *Philosophy and Grammar*, Amsterdam: Reidel, pp. 79–100.

Lewis, D.: 1996, Elusive knowledge, *Australasian Journal of Philosophy* **74**(4), 549–567.

MacFarlane, J.: 2014, *Assessment sensitivity: Relative truth and its applications*, OUP Oxford.

Montague, R.: 1960, Logical necessity, physical necessity, ethics, and quantifiers, *Inquiry* **3**(1), 259–269.

Montague, R.: 1968, Pragmatics, *in* R. Klibansky (ed.), *Contemporary Philosophy: A Survey*, Vol. 1, Florence, La Nuova Italia Editrice, pp. 102–22. Reprinted in *Formal philosophy: selected papers of Richard Montague*, 1974.

Montague, R.: 1970a, Pragmatics and intensional logic, *Synthese* **22**(1), 68–94.

Montague, R.: 1970b, Universal grammar, *Theoria* **36**(3), 373–398.

Ninan, D.: 2010, Semantics and the objects of assertion, *Linguistics and Philosophy* **33**(5), 335–380.

Prior, A.: 1956, Modality and quantification in S5, *The journal of symbolic logic* **21**(1), 60–62.

Prior, A.: 1968a, Egocentric logic, *Nous* **2**(3), 191–207.

Prior, A.: 1968b, Now, *Nous* **2**(2), 101–119.

Rabern, B.: 2012, Against the identification of assertoric content with compositional value, *Synthese* **189**(1), 75–96.

Rabern, B.: 2013, Monsters in Kaplan's logic of demonstratives, *Philosophical Studies* **164**(2), 393–404.

Rabern, B.: 2016, The history of the use of ⟦.⟧-notation in natural language semantics, *Semantics and Pragmatics* **9**(12).

Schaffer, J.: 2004, From contextualism to contrastivism, *Philosophical Studies* **119**(1), 73–103.

Schaffer, J.: 2015, Lewis on knowledge ascriptions, *A Companion to David Lewis* **57**, 473.

Scott, D.: 1970, Advice on modal logic, *in* K. Lambert (ed.), *Philosophical Problems in Logic: Some Recent Developments*, D. Reidel, pp. 143–173.

Stalnaker, R.: forthcoming, David lewis on context, *in* D. Ball and B. Rabern (eds), *The Science of Meaning*, Oxford : Oxford University Press.

Stanley, J.: 1997b, Names and rigid designation, *in* B. Hale and C. Wright (eds), *A Companion to the Philosophy of Language*, Oxford, Blackwell Press, pp. 555–585.

Stanley, J.: 2002, Modality and what is said, *Noûs* **36**(s16), 321–344.

Stanley, J.: 2005, *Knowledge and Practical Interests*, Oxford University Press.

Vlach, F.: 1973, *'Now' and 'Then': a formal study in the logic of tense anaphora*, PhD thesis, UCLA.

Von Fintel, K. and Heim, I.: 2011, *Lecture notes on intensional semantics*, MIT.

Williamson, T.: 2002, *Knowledge and its Limits*, Oxford University Press.

Wright, C.: 2004, Warrant for nothing (and foundations for free)?, *Aristotelian Society Supplementary Volume*, Vol. 78, pp. 167–212.

Yalcin, S.: 2007, Epistemic modals, *Mind* **116**(464), 983–1026.

Yalcin, S.: 2014, Semantics and metasemantics in the context of generative grammar, *in* A. Burgess and B. Sherman (eds), *Metasemantics: New Essays on the Foundations of Meaning*, Oxford University Press.