Contextualism preserved

James Ravi Kirkpatrick
Somerville College, University of Oxford, UK
james.kirkpatrick@philosophy.ox.ac.uk

This paper aims to reconcile the context-sensitivity of natural language with the essential epistemological role that language plays in the preservation and transmission of content.

1 Introduction

There is an intuitive and compelling view about the epistemological significance of language to our privates lives as reasoners and our social lives as communicators. According to this view, language plays an essential role in giving expression to the information that we gather and later store in memory. Once stored, this information survives robustly across time to be recalled at some later date for use in theoretical reasoning or in practical deliberation about means to desirable ends. We can also communicate this information to others through language, sharing our beliefs and intentions, so that we can coordinate and draw upon our shared information in the pursuit of joint projects and mutual goals. In a slogan, language serves as a vehicle for the preservation and transmission of information.

However, this picture is threatened by the pervasive context-sensitivity of natural language.[1] For it is widely held in the philosophy of language and linguistics that the semantic contributions of many expressions in natural language are sensitive to the contexts in which they are uttered. Some kinds of context-sensitivity are familiar, such as the context-sensitivity involved with indexicals and demonstratives like *I*, *here*, *now*, *this* and *that*. For example, when I say the sentence 'I am happy', I express something different from when you say 'I am happy', since the meaning of the first-personal pronoun *I* means something different in my mouth to what it means in yours. But over the last fifty years theorists have argued for increasingly more complicated forms of context-sensitivity, extending the range of expressions purported to be context-sensitive to include, quantifiers, gradable adjectives, propositional attitudes, modals, conditionals, and many more.

[1] See, for example, Hawthorne (2004); Williamson (2005a,b, 2020).

While the postulation of context-sensitivity is ubiquitous in our linguistic and philosophical theorising about natural language, its existence would seem to frustrate the central role that language plays in the preservation and transmission of content. For if language is to play this role, then our words and sentences cannot just be one-off devices, to be cast aside after use. Rather, the meanings of words and sentences must, in some sense, be fundamentally stable across contexts. But it is unclear how to assimilate these roles with widespread context-sensitivity. For if the preservation and transmission of content crucially involves the preservation and transmission of *context-sensitive* content, then the contents of our beliefs and desires would ebb and flow with the shifting sands of context. At the extreme, widespread context-sensitivity would critically inhibit the preservation and transmission of information with disastrous consequences to theoretical and practical reasoning.

In this paper, I aim to reconcile the context-sensitivity of natural language with the role that language plays in the preservation and transmission of content. The point is not to argue for the context-sensitivity of any particular expression nor to argue for any particular theory of context-sensitivity. Rather, I aim to disarm a certain kind of objection to context-sensitivity that takes it as incompatible with the centrality of language in preserving information in memory and transmitting information in communication. The crucial part of the strategy will be to articulate a theory of semantic content and information storage that is compatible with context-sensitivity and which does not do violence to memory and communication.

This paper is structured as follows. Section 2 outlines the varieties of contextualism in natural language. Section 3 recounts how context-sensitivity poses a problem for the epistemological roles which language is meant to play, paying particular attention to the presentations of this argument made by John Hawthorne (2004) and Timothy Williamson (2005a; 2005b; 2020). Section 4 raises issues with the picture of belief at the centre of Hawthorne's and Williamson's arguments and argues that it should be rejected. Section 5 advocates for an alternative picture of belief, which draws on the work of Robert Stalnaker, David Lewis, and David Kaplan, and argues that this picture dissolves the problem of content preservation and transmission for contextualism.

2 Varieties of Contextualism

In its broadest formulation, contextualism about linguistic meaning is the view that the meanings of certain classes of sentences are context-sensitive. In particular, such sentences are context-sensitive in virtue of containing certain expressions that are context-sensitive. Contextualism is typically defended by appealing to our intuitions about 'what is said' by utterances of sentences in different contexts and their truth values. For example, if someone in one context can truly say a sentence, while someone else in a different context can simultaneously use the same sentence and say something false, this suggests that the sentence in question is context-sensitive. Theorists have used this strategy to defend contextualism across a wide variety of linguistic expressions, such as indexicals, demonstratives,

gradable adjectives, quantifiers, propositional attitudes, modals, and conditionals, among others.

Most theorists agree that contextualism is uncontroversially true for at least some of these expressions.[2] Consider, for example, indexicals, such as *I*, *here* and *now*, and demonstratives, such as *this* and *that*. For such expressions, the speaker's context, together with a rule (perhaps together with a gesture), determines the meaning of that expression in that context. If I say 'I am here now', I express something different from when you say 'I am here now'; we automatically apply the rule that *I* as used in a given context refers to the speaker in that context, and so when I use *I*, I refer to me, and when you use *I*, you refer to you. Similar rules map *here* as used in a given context to the location of the speaker in that context and *now* as used in a given context to the time when the utterance was made. Analogous arguments show that demonstratives are context-sensitive. When I say 'That is a dog' and you say 'That is not a dog', we may do not necessarily disagree, even though your sentence is the verbal negation of mine. After all, a demonstrative refers to that which the speaker who uses it intends to refers to by using it, and we may be referring to different things.

Other kinds of context-sensitivity are more controversial and the process of determining the contextual meanings of such expressions is more complicated. Consider, for example, quantifiers (e.g., *every*, *some*) and gradable adjectives (e.g., *tall*, *flat*).[3] When I say 'Every beer is in the fridge' and you say 'Every beer is not in the fridge', we may not be in disagreement, since I could mean that every beer in my house is in the fridge, while you might mean that every beer in the supermarket is not in the fridge. That is, the domain over which our quantifiers range may vary from context to context. Similarly, when I say that 'John is tall', and you say 'John is not tall', speaking of the same person, we may not be in disagreement, since I might be in a conversation about jockeys, and you might be in a conversation about basketball players.[4] That is, what I have said might be true relative to the standards of tallness in my context, and what you said might be true relative to different standards in your context. If someone else wanted to report our utterances, she had better not stick to a single context and merely repeat our words; rather she should make clear which beers we are talking about and what standards for tallness are salient in our respective contexts. After all, the beers that I was talking about were a proper subset of the beers you were talking about. And when you said 'John is tall', the word *tall* seemed to express the property of being tall for a basketball player, whereas when I said 'John is tall', the word *tall* seemed to express the property of being tall for a jockey. Similar remarks apply to other purportedly context-sensitive expressions mentioned above.[5]

[2] For discussion of obvious and non-controversial instances of linguistic context-sensitivity, see Borg (2004, 2012) and Cappelen and Lepore (2005: ch. 7).

[3] For a defence of contextualism about quantifiers and gradable adjectives, see, e.g., Stanley and Szabó (2000) and Kennedy (1999, 2007), respectively.

[4] Context-sensitivity should not be confused with vagueness, even though many context-sensitive expressions are vague; see Williamson (2005b: 99).

[5] For discussion of the context-sensitivity of conditionals, see, e.g., Stalnaker (1968); Lewis (1973);

There are many ways that contextualism can be implemented. For example, it may be that the expression itself is context-sensitive, and so its semantic contribution varies from one context to another. Or it may also be that there are covert variables or parameters that in some way associated with an expression, and the value of these variables and parameters may be dependent on the context. These variables and parameters may syntactically associated with the expression, but they may also be associated only once the sentence has been enriched through some pragmatic process or at in the level of Language of Thought.[6]

Discussing all these possible views is outside the scope of this paper. Instead, I will focus specifically on an implementation of contextualism within formal semantics which holds both (a) that semantic content is compositional (that is, the meaning of any complex expression is constituted by the meanings of its parts, together with syntactic rules) and (b) that any (relevant) context-sensitivity is either encoded in the meaning of the expression itself or syntactically associated phonologically null constituents. The motivation for imposing these constraints is not to deny that linguistic context-sensitivity can take any other form, but rather to focus our sights on a well-behaved form of contextualism. For if no solution to the problem of preservation and transmission can be found in this framework, there is little hope for more complex contextualist implementations.[7]

3 The Problem of Preservation and Transmission

Contextualism is typically defended by appealing to our intuitions about the truth-values of sentences in different contexts with little consideration for how context-sensitive information may be retained and transmitted across contexts. But the existence of context-sensitive information seems to be at odds with how we use language to communicate about the world, as well as how we reason and decide how to act. We rely on the wide variety of information communicated by others in deciding what to do; we have the ability to communicate our mental states and represent the mental states of others on this basis, as well as the ability to adapt previously accepted information to novel contexts. But this ability to adapt old information to novel contexts lies in tension with the idea that context-sensitive expressions mean different things across different contexts. This tension is given explicit and highly illuminating discussion by John Hawthorne (2004) and Timothy Williamson (2005a; 2005b; 2020).[8] In this section, I shall set out some salient features

for modals, see, e.g., Egan and Weatherson (2011); Kratzer (2012); and for knowledge ascriptions, see, e.g., Lewis (1996); Cohen (1999).

[6] For a good discussion of the variety of forms of contextualism, at least with regards to quantifier domain restriction, see Elbourne (2008).

[7] For other more radical contextualist theories of meaning, see (e.g., Sperber and Wilson 1986; Carston 2002; Travis 1996; Searle 1980).

[8] See especially, Hawthorne (2004: 108–111) and Williamson (2005b: 99–102). I want to stress that Hawthorne and Williamson are concerned mainly with arguing against contextualism about belief ascriptions and knowledge ascriptions, and conditional constructions, although they suggest that their arguments extend to "other philosophically contentious contextualist claims", the context-

of Hawthorne's and Williamson's discussions, with a view to raising the question of whether the tension they describe undermines contextualist theories of meaning.

To see this tension, let us focus on belief and consider a simple model of belief according to which we all have a 'belief box' that contains sentences in English.[9] According to this model, to believe a proposition p is to have a sentence in one's belief box that expresses the proposition that p. It is easy to see that this model comfortably handles beliefs about context-insensitive contents, such as the belief that $1 + 1 = 2$. For according to this model, one believes the proposition that $1 + 1 = 2$ just in case one has a sentence in one's belief box that expresses that proposition, say, for example, the following sentence:

(1) $1 + 1 = 2$.

Mathematical terms have context-invariant contents, if any do, so (1) expresses the proposition that $1 + 1 = 2$ regardless of the context of utterance. Then, the simple model of belief delivers a simple and elegant explanation of how we preserve and transfer such content across contexts: given that (1) expresses the same proposition in every context, we can safely store (1) in our belief box with the knowledge that we can retrieve the same content from it in any context, and so we can also transmit it that content to any competent speaker of English simply by uttering (1).

However, things get tricky once we introduce context-sensitivity into the picture. Suppose that a speaker notices that it is raining and says the following sentence:

(2) It is raining here.

Sentences like (2) express different propositions in different contexts of utterance. For example, when (2) is uttered in Oxford, it expresses the proposition that it is raining in Oxford, and when it is uttered in Cambridge, it expresses the proposition that it is raining in Cambridge. Consequently, if the speaker wishes to preserve the information that (2) expresses in its original context of utterance, it will not do to store (2) as a permanent inhabit in the belief box, since the sentence will change its content as the speaker moves through time and space. That is, as the speaker changes location from Oxford to Cambridge, she will no longer be able to access

sensitivity of which "is not manifest to ordinary speakers" (Williamson 2005b: 102). Nevertheless, I do not wish to suggest that Hawthorne and Williamson would endorse all the arguments in this section.

[9] While this toy model may seem overly simplistic, it is essentially a representational theory of belief and more sophisticated analogues of this model are widely accepted among theorists. For example, rather than storing sentences in English, we can imagine that the belief box stores mental representations that are sentences in an internal language of thought; call this language 'mentalise' (cf. Fodor 1975, 2008). According to the language of thought hypothesis, for each natural language sentence, there is a corresponding sentence in the language of thought that is a representation of the same information as that natural language sentence. On this view, an agent believes a proposition just in case she has a representation of that proposition that plays the right kind of role in her cognition, namely, it is deployable in the ways that are characteristic of belief, such as in theoretical inferences or practical deliberations. See Section 4, for further discussion.

the old information that it was raining in Oxford by retrieving (2) from her belief box, since the content that it expresses will have changed along with her location.

How can the speaker preserve in her memory the information that she got from the original utterance of (2)? A natural suggestion is that she might try preserving the information in her belief box with a non-shifty sentence that, in new contexts, expresses the same information that (2) did in the original context. For example, she might try rigidifying the sentence, by replacing 'here' with a description of the time and place of utterance, or perhaps storing the sentence with a mental snapshot or file of the relevant spatiotemporal location, like in the following:

(3) It was raining in Oxford at noon on December 12th.

(4) It was raining there [accompanied by a mental snapshot or file of the relevant spatiotemporal location].

However, such strategies do not help when one does not know the value of the contextual parameters that affect the meaning of the sentences (Williamson 2005b: 100–101).[10] If you do not know where you are or what time it is, you would not know how to replace 'here' and 'now' with a context-insensitive designation of the time or place. And existentially quantifying out the contextual parameters with 'somewhere' and 'sometime' will likely lead to substantial information loss, thus failing to preserve the same information as the original sentence expressed. While such observations illustrate the utility of indexicals, they also highlight their limitations in preserving and transmitting information.[11]

This problem is pervasive and things are even trickier for those expressions whose context-sensitive components vary with more elusive parameters. Suppose that I say (5) and at the time of utterance I express the same property as you do by 'tall':

(5) John is tall.

Suppose you accept what I say and, using the tricks from above, you store something like the following in your belief:

(6) John is tall [plus a date index or a mental snapshot of the relevant context].

[10] For related discussion of how contextualism predicts that we are 'blind' to the semantic workings of our language, see Schiffer (1992, 1994, 1995).

[11] Another strategy would be that if in c^* you remember that you acquired S in c, and you remember what c was like, you might be able to construct another sentence, S^*, that expresses the same truth in c^* as S expressed in c. However, as Williamson (2005b) argues that this strategy imposes a heavy burden on memory. While most people cannot remember where they got much of the information on which they rely, this does not mean that they are not warranted in relying on that information to guide action or to pass on to others. If one cannot remember from where one got one's information, then it is unclear how one could reconstruct another sentence that now expresses the same information as the original sentence did in the original context.

But now suppose that your standards for tallness rise as the contextually salient comparison class changes. Your belief about John's height may now come out false — as will many of your other once true beliefs — unless you somehow update the sentences in your belief box as and when the standards for tallness shift. One might try to store the contextually salient comparison class; conceivably (5) as uttered when jockeys are the salient comparison class expresses the same proposition as the sentence in (7):

(7) John is tall for a jockey.

But the average height for jockeys also varies across times and places, and so without a method for fixing context-insensitive content, it is unclear what eternal context-insensitive sentence invariably expresses the same information in every context as the original sentence does in the original context of utterance. And in any case, if you don't know which comparison class was salient in my context, then what are you meant to store? One could try a description of the context, but such descriptions will be of little use to those who do not know what was going on at that place and that time. Similar problems arise for expressions involving other subtle forms of context-sensitivity, including propositional attitudes, modals and conditionals, and the more extensively a given expression varies with context, the worse the problems with preserving and transmitting information will be.

In more general terms, suppose that in a context c, you acquire an item of information expressed by sentence S as uttered in c. Suppose you then stored this information by storing S in your memory. But when the time comes to retrieve that information, for guiding action or passing it on to someone else, you will be in a new context c^*, and even if S expressed a truth in c, there is no guarantee it will express a truth in c^*. Storing S along with an explicit index or mental snapshot does not help nor does trying to reconstruct a new sentence S^* that expresses the same content in c^* as S did in c. If we store information by storing these sentences in our belief box, then if the context changes, there is no guarantee that we will have cognitive access to those true propositions in our belief box. If we transmit information across contexts by recalling and repeating the sentences in our belief box, there is no guarantee that the information transmitted is the same as the information first stored: "you will no longer have cognitive hold on those true propositions in your belief box once truly expressed that, as a result, you once truly believed" (Hawthorne 2004: 110). And we shouldn't hold out hope for the existence of eternal sentences that invariantly expresses the same information as shifty sentences do on particularly occasions of use; as as Williamson writes, "if such an eternal sentence [...] were readily available for the occasional sentence [...] on each occasion, then the context-relativity of [the occasional sentence] would be, if not idle, at least underemployed" (Williamson 2005b: 100).

The gauntlet has been thrown. Any attempt to reconcile the pervasiveness of context-sensitivity in natural language with the essential role that language plays in the preservation and transmission of information must provide an explanation of how we preserve the information expressed by context-sensitive sentences that

in some context-insensitive way, as well as how these mechanisms work when we are ignorant about features of the original context of utterance.[12]

4 The Linguistic Picture of Belief

I want to suggest that the problem of content preservation and transmission arises from the specific conception of belief upon which Hawthorne and Williamson rely in their arguments. This conception of belief — one that is essentially linguistic in something like the sense sketched in the previous section — is criticised by Robert Stalnaker in *Inquiry*, in the context of a wider discussion of the problem of intentionality, that is, the problem of explaining the nature and possibility of minds and mental states to represent.[13] In this section, I shall set out some salient features of Stalnaker's discussion, with a view towards undermining any appeal to the linguistic picture in setting up the problems of content preservation and transmission. In particular, I want to suggest that the linguistic picture is committed to far more radical conclusions than are tolerable, which should motivate us towards an alternative picture of belief that better accommodates the context-sensitivity of natural language.

Let us take for granted that what it is to have a belief or desire — and, more generally, an intentional or representational mental state — is to stand in an appropriate relation to a proposition. According to the linguistic conception of belief, what it is to stand in the appropriate relation to a proposition is to stand in an associated relation to some sentence — perhaps a sentence of natural language, perhaps a sentence in a 'language of thought' — where the sentence expresses or 'means' the proposition.[14] Thus, the linguistic picture attempts to solve the problem of intentionality "by reducing intentional relations to semantic relations and then giving a naturalistic explanation of semantic relations" (Stalnaker 1984: 27). In doing so, the analysis divides the problem of providing a naturalistic account of belief into two subproblems. The first problem is to explain what it is to stand in an associated relation to some sentence, such as 'storage' in a 'belief box' or, less metaphorically, in terms of dispositions to assent to the sentence or to use the sentence in question in reasoning and deliberation. The second problem is to explain what it is for a sentence to mean that proposition, which is more difficult to explain in naturalistically acceptable terms. The belief box model is a variant of this linguistic picture of belief, since it holds that one believes a proposition just in case there is a sentence stored in one's belief box that expresses that proposition. Thus, by undermining the viability of the linguistic picture of belief, we in turn

[12] Other responses to this argument might deny that information is preserved exactly across contexts, but so long as what is preserved is similar enough to what was originally expressed for practical purposes, there is little epistemologically peril; see, for example, Bezuidenhout (2002: 212-213) and Sperber and Wilson (1986: 192).

[13] See especially Chapters 1, 2, and 4.

[14] See Field (1978) for a clear and explicit development of this idea. Stalnaker's discussion focuses largely on Field's proposal, which he takes to be representative of the linguistic picture of belief.

undermining the viability of the belief box model.

Stalnaker considers two versions of the linguistic picture — an atomistic version as advocated by Hartry Field and a holistic version drawn from the writings of Donald Davidson — and argues that neither provides an adequate naturalistic solution to the problem of intentionality.

The atomistic version attempts to give a naturalistic explanation of what it is for a sentence to mean some proposition in terms of a Tarskian truth theory supplemented with a causal theory of primitive denotation.[15] Stalnaker draws upon Field's own observation (see Field (1972)) that Tarski's treatment of atomic terms is inadequate as an explanation of primitive denotation to make a parallel point can be made about a Tarskian treatment of constructions and operations. That is, a Tarskian truth theory merely enumerates what specific atomic terms denote, without providing an account that predictively generalises over unlisted expressions nor an explanation for why such expressions are mapped to their denotations. Similarly, the recursive clauses in a Tarskian truth definition merely enumerates what specific functional expressions contribute to the truth conditions of sentences in which they occur, establishing "an extensional equivalence between ascriptions of truth to sentences and ascriptions of primitive denotation to the ultimate constituents of the sentences, but [without telling] us any more about the basis for those ascriptions than does a definition of primitive denotation by enumeration", and so they cannot be generalised to new recursive expressions (Stalnaker 1984: 30). Since Field's attempt to reduce truth to primitive denotation relies on these recursive clauses, he must provide an account of how the relevant kinds of expressions mean what they do, which is something that a mere enumeration of what they individually mean cannot do: while a Tarskian truth theory can explain "how a sentence comes to have the truth conditions which it has", it cannot explain "what physical (or sociological or psychological) facts constitute the fact that sentence S expresses proposition x" (Stalnaker 1984: 32).

Stalnaker goes on to criticise Field's claim that primitive denotation can be explained in terms of a non-intentional causal relation that maps atomic expressions to suitable denotations. As Stalnaker points out, the kinds of causal connections evoked in causal theories of reference are typically described in intentional terms. For example, in *Naming and Necessity*, Saul Kripke characterises the causal connection involved in his theory of reference partly in terms of speakers' intention to preserve reference.[16] We have little reason to believe that the relevant causal connections between expressions and their denotations can be given in non-intentional terms. Thus, it is doubtful that the atomistic version can provide a naturalistic account of linguistic meaning that makes no reference to mental representation.

The holistic version of the linguistic picture that Stalnaker considers (see Stalnaker 1984: 36–41) is drawn from the work of Donald Davidson.[17] According to

[15] See, in particular, Stalnaker (1984: 28–33).

[16] "Notice that the preceding outline hardly *eliminates* the notion of reference; on the contrary, it takes the notion of intending to use the same reference as given" (Kripke 1972: 302).

[17] Stalnaker is careful to note that Davidson himself does not attempt to explain semantic notions in non-intentional terms. "My way of trying to give an account of language and meaning makes

this view, the semantic relation between a sentence and its truth conditions is taken to be more primitive than the semantic relation of primitive denotation. Stalnaker labels this holistic linguistic approach a "theory of interpretation" and describes an oversimplification of it as follows:

> The main task of the theory of interpretation is to give a systematic procedure for dividing the sentences [of the language of a speech community] into two classes, the division to be depend on the way the world is. On the basis of a small finite number of postulates establishing correlations between the basic sound or shape types and things and properties, together with recursive rules classifying complex utterance types on the basis of these correlations, the theory defines a class of *true* sentences. The empirical claim of the theory is that the actual utterances will normally be drawn from this class. (Stalnaker 1984: 36–37)

Such a theory would provide a naturalistic account of meaning, but at the cost identifying and explaining linguistic action in isolation from other kinds of action. For example, speech is generalised recognised as such within its role in a wider range of rational activity, the latter of which must be given in intentional terms. Any theory of interpretation in non-intentional terms would sharply divide speech from rational action more generally. Furthermore, belief plays a characteristic role in the explanation of non-linguistic as linguistic behaviour, such as in the explanation of the actions of non-linguistic animals. While non-linguistic animals may not have certain kinds of beliefs, such as metalinguistic beliefs about representational content, the theory of interpretation would preclude any explanation of their behaviour in terms of beliefs (cf. Stalnaker 1984: 62–64). Such a theory has little plausibility.

So far, we have seen that Stalnaker's objections to the linguistic picture of belief, while not conclusive, give reason to doubt that the general adequacy of the picture of belief that Hawthorne and Williamson draw as a naturalistic theory of belief in a wider discussion of the problem of intentionality. However, for present purposes, the most important criticism that Stalnaker raises against the linguistic picture of belief directly concerns the compatibility of the linguistic model of belief with context-sensitivity. In particular, Stalnaker calls into question the very idea of an autonomous semantic theory — that is, a semantic theory that assigns content to sentences independently of the attitudes of speakers — needed to make the linguistic picture of belief work. For such a theory cannot accommodate certain context-sensitive expressions, such as demonstratives or quantifiers, on pain of vicious circularity, since the denotations of demonstratives or quantificational domains are plausibly determined in part by intentional features of the context,

essential use of such concepts as those of belief and intention, and I do not believe it is possible to reduce these notions to anything more scientific or behaviouristic" (Davidson 1973: 4–5). Stalnaker's objections concern only those theories of interpretation that appealed to in a naturalised account of mental representation.

namely, the attitudes and intentions of speakers. The underlying semantic theory of that language cannot be given in non-intentional terms, and so the linguistic picture of belief for communities who speak such languages cannot provide a non-circular account of belief.

Once we admit that natural language contains at least some context-sensitive expressions the denotations of which must be resolved with reference to speakers' intentions and beliefs, Hawthorne and Williamson face a dilemma. The first horn of the dilemma is that they admit that the meaning of some context-sensitive expressions are functions of speakers' intentions, and thus face the challenge of explaining how the linguistic theory of belief that they presuppose is capable of providing a naturalistic account of belief. As I have argued in this section, Stalnaker provides compelling reasons against taking this horn. The second horn is that they deny that a non-autonomous semantics plays any role in a general theory of belief. But this comes at the cost of denying that demonstratives and other referring expressions are context-sensitive expressions, a cost which most theorists should consider too great. For even semantic minimalists — theorists such as Emma Borg, Herman Cappelen, and Ernie Lepore who accept only the existence of an extremely limited number of context-sensitive expressions — count demonstratives among the context-sensitive expressions of natural language.

One may be tempted to retain non-autonomous semantics for natural language and move to context-independent inner language of thought, the terms of which a naturalised linguistic theory of belief would be given. But to make this move, Hawthorne and Williamson would be affording themselves the very conceptual resources that they denied the contextualist in explaining away the problem of content preservation and transmission, namely, a context-independent language that serves as an invariant medium to preserve the content of a context-sensitive natural language language. The burden of proof lies on them to explain why the resources of this context-independent language of thought cannot then be extended to the more contentious forms of context-sensitivity with which they take issue. Furthermore, as Stalnaker points out, adopting this strategy in effect denies that intentionality is essentially a social phenomenon: "If this is the motivation for holding that language is essential to thought, then the language must be a language actually used by members of communities of speakers" (Stalnaker 1984: 41).

Hawthorne's and Williamson's argument against contextualism either proves too much or too little. By appealing the linguistic picture of belief, their arguments either shows that natural language is not context-sensitive or they fail to demarcate between those forms of context-sensitivity that fall foul to their argument and those which do not. This, in turn, gives us reason to suspect that the culprit is not context-sensitivity, but rather the picture of belief on the objection is premised. I submit, then, that it is the linguistic picture of belief that must go.

5 The proposal

The previous section raised issues with the linguistic picture of belief on which Hawthorne's and Williamson's objections are premised. But we do not need sentence-like entities to play the role of the middleman between semantic content expressed and beliefs had. Instead, we can store the semantic content that context-sensitive sentences express at a context directly, with the hope that anything explained in terms of standing in an associated relation to such sentence can be explained better in terms of standing in relation to the content itself. The goal of this section is to argue that, by cutting out the middleman and modelling beliefs directly as propositions, the problem of content preservation and transmission does not get off the ground. To be clear, the core components of the proposal are not novel; they are drawn from familiar work by Stalnaker, Lewis, and Kaplan. However, as far as I am aware, their application to the problem of content preservation and transmission is novel. I will first outline the proposal in informal terms (Section 5.1) and then make the proposal more precise (Section 5.2). I then demonstrate how the resulting framework handles the preservation of content in a context-sensitive language, both when one knows the relevant features of the context (Sections 5.3 and 5.4) and when one does not know the relevant features of the context (Section 5.6)

5.1 Basics

Let us take the semantic content expressed by a sentence at a context to be a proposition, modelled here as a function from possible worlds to truth-values. Then we may model the content of a belief as the set of worlds at which that belief is true. Moreover, we can model the entirety of what one believes, one's *belief state*, as a set of possible worlds, namely, the possible worlds that are compatible with what one believes. The effect of updating one's beliefs, on this view, is to eliminate possibilities that are incompatible with the proposition asserted from one's information state. As our knowledge grows, we exclude worlds from our belief set, gradually narrowing down which possibilities we think might be actual.[18]

To demonstrate how this picture of belief dissolves the problem of content preservation and transmission, let me explain how a compositional semantics for natural language delivers propositions to serve as the objects of belief. According to a compositional theory of meaning, the meaning of a complex expression is determined by the meanings of its constituent expressions and the rules used to combine them. In extensional semantics, semantic values are extensions: sentences

[18] The limitations of this model are well-known; the usual cavaets and qualifications apply. Hyperintensional contents that should be distinguished are no longer distinguished; logical truths are true at every possible world, so every logical truths is believed. Logical equivalences are believed if one believes any one equivalent, since logical equivalents are truth at exactly the same worlds. Any logical entailment of any proposition which one believes is also believed. Such limitations are acceptable for present purposes; from the point of view of abstraction, they do no damage to the conclusions we wish to make.

denote truth-values; names denote individuals; intransitive nouns and verbs denote subsets of the domain of discourse (or their characteristic functions); and so on. Atomic sentences are formed in accordance with the syntax or grammar of the language and corresponding composition rules determine their denotations from the denotations of its constituent expressions. Compound sentences are then built from sentences together with truth-functional connectives, so that the truth-value denoted by a compound sentence will be a function of the truth-values of its constituent sentences.

However, some expressions, such as modal expressions, are not truth-functional, and so a compositional semantics for a modal language requires a more fine-grained notion of semantic value. Possible world semantics is an intensional semantics that provides such a notion of semantic value by generalising extensional semantics. Broadly speaking, an intension is a function from possible worlds to extensions. If the extension of a name is an individual drawn from the domain of discourse, the intension of the name will be an individual concept, that is, a function from worlds to individuals. If the extension of a sentence is a truth-value, the intension of such a sentence will be proposition, that is, a function from worlds to truth-values. If the extension of an intransitive noun or verb is a subset of the domain of discourse (or its characteristic function), the intension of such a predicate will be a function from individuals to functions from worlds to truth-values. Extensional operators and semantic composition rules generalise to their intensional analogues. Truth becomes a relative notion; a sentence will no longer be true simpliciter, but rather true relative to a world. And so on. The upshot of this conception of semantic value is that a range of intensional operators, such as modal expressions like *may* and *must*, can now be interpreted as functions that operate on predicate and sentence denotations in a non-extensional manner. The resulting system allows for the interpretation of modal language and provides a notion of informational content that allows us to express a thought with a certain content.

A further generalisation to intensional semantics allows us to accommodate context-sensitive expressions. Just as modal expressions in an intensional semantics are sensitive to a possible world, context-sensitive expressions are sensitive to certain features of the context, such as the speaker, the audience, the time of utterance, the place of utterance, and so on. Thus, the notion of an intension must be generalised further by incorporating the notion of an *index* (Montague 1968, 1970). That is, we must generalise the notion of an intension to a function from indices to extensions, where an index is an ordered tuple of items that includes anything on which the extension of an expression might depend. For technical reasons, it will be necessary to divide the Montagovian index into two parts: a *context*, which is something like Montague's index (that is, a tuple of items determined by the concrete situation in which the utterance takes place and which is held fixed across compositional interpretation), and a *circumstance*, which for present purposes

we may take to be a world (Kaplan 1989).[19,20] Then, the compositional process of interpreting a sentence goes by two steps: we first take linguistic meaning of an expression — call this its *Kaplanian character* — and use it, together with the context, to determine its semantic value — call this its *Kaplanian content*.[21] For example, the character of a sentence is a function from contexts to propositions, that is, functions from indices to functions from circumstances to truth values; the content of a sentence in a context is the proposition expressed by that sentence in the context as represented in Figure 1. The character of the first-personal pronoun is a function from contexts to individuals; the content of a first personal pronoun at a context of utterance is the speaker of the utterance. And so on.

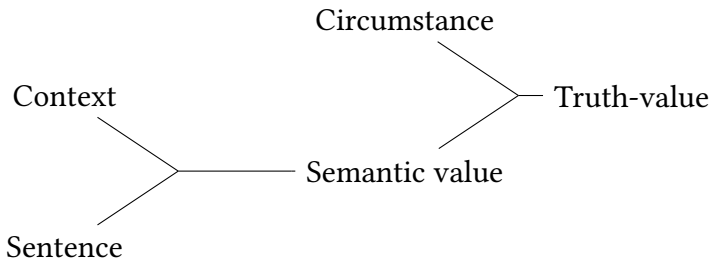

Figure 1: The Kaplan–Lewis model of semantic content

I want to suggest that if natural language involves pervasive context-sensitivity, then the very possibility of content preservation and transmission requires that what we store in our beliefs involves something like the notion os semantic content as given under this model of semantics.

5.2 Information states and updates

To see how this model of belief and content resolves the problem of content preservation and transmission, it will be useful to sketch a toy model for how our belief sets as we learn new information, as well as a toy semantics for a context-sensitive fragment of natural language. There are various ways of implementing this picture. I elect to approach the issue from the perspective of a relatively 'textbook' model

[19] In Kaplan's original theory, the index consists of a distinguished context, the elements of which can shift or change within the scope of some operator. However, assuming that modals, tenses, locative expressions are analysed in terms of covert object language quantifiers ranging over covert world, time, and location variables, we can do away with most of this machinery; cf. Cresswell (1990); King (2003); Schaffer (2018).

[20] If we wish to model indexical beliefs, we may want even more fine-grained indices, such as *centered worlds* (tuples of individuals at a world and time) and more fine-grained contents, such as *centered propositons* (functions from centered worlds to truth-values); see, for example, Lewis (1979a); Perry (1977, 1979), although, see Cappelen and Dever (2013); Magidor (2015) for dissenting opinions.

[21] David Lewis (1980) argues against the two-step approach to semantic interpretation, instead taking meaning to be function from a context–index pair to a truth-value. For illuminating responses to Lewis's argument, see King (2003), Stalnaker (2014: 22ff.), and Schaffer (2018).

of belief updating in a dynamic semantics and relatively 'textbook' semantic analyses for context-sensitive expressions. In theory, nothing hinges on the particular semantic analyses I assume; advocates of more exotic analyses should substitute their preferred views as and when.

To model the evolution of beliefs in response to new information, we need two things: a definition of *information states* (I use this term interchangeably with 'belief state') and an *update function* that provides an instruction for each sentence in our language on how to change an information state in response to the informational content of that sentence.[22]

Definition 5.1. A possible world w assigns a truth value to atomic sentence α at a context c. $\mathcal{W}$ is the set of all possible worlds. An information state s is a set of possible worlds. The set of information states $\mathcal{S} = \wp(\mathcal{W})$. An update function at a context c, $[.]^c$ is a function from sentences to functions from information states to information states. $s[\phi]^c$ is the result of inputting s into $[\phi]^c$.

Once we have a representation of information states, we can recursively define our update function, $[.]^c$. Intuitively, atomic sentences narrow down an information state by eliminating the worlds in which they are false. Negations remove any worlds that would survive updating with its prejacent. A conjunction first updates a state with the first conjunct, then updates the result with the second conjunction:[23]

Definition 5.2. If $[.]^c$ is an update function at a context, c is a context, α is an atomic sentence, and ϕ, ψ are sentences, then:

$$\begin{aligned} s[\alpha]^c &= s \cap [\![\alpha]\!]^c \\ s[\neg\phi]^c &= s - s[\phi]^c \\ s[\phi \wedge \psi]^c &= s[\phi]^c[\psi]^c \end{aligned}$$

We can say that an information state supports a sentence ϕ just in case updating the state with ϕ has no effect.

Definition 5.3. s SUPPORTS ϕ in c iff $s[\phi]^c = s$

In contrast with the linguistic picture of belief, the acquisition of new beliefs is not modelled in terms of storing sentences, but in terms of narrowing of our belief set by excluding worlds that are incompatible with the newly acquired information. As we acquire new beliefs via testimony, we exclude those worlds in our belief set that are incompatible with the semantic content of testimonial sentences in the context in which they are uttered. This information is preserved across contexts, since once our belief states are suitably updated, they remain consistent with the semantic content of those testimonial sentences in their utterance context.

[22] For seminal papers in the development of these resources, see Stalnaker (1973); Heim (1982, 1983); Veltman (1985, 1996); Groenendijk and Stokhof (1991), among others.

[23] '$[\![.]\!]^{c,w}$ denotes the interpretation function of the model of the language, which maps well-formed expressions to their extensions relative to context c and a possible world w.

5.3 Context-sensitivity under certainty

To get a sense of how this framework works, let us consider a context-insensitive fragment of English. Consider what happens when we update an information state s with the information expressed in (8):

(8) Mary smokes.

I shall assume that proper names denote individual concepts (functions from worlds to individuals) and intransitive verbs denote intensionalised properties (functions from individual concepts to functions from contexts to propositions). Then the lexical entries for *Mary* and *smokes* are as follows:

(9) $[\![\text{Mary}]\!] = \lambda c.\text{Mary}$

(10) $[\![\text{smokes}]\!] = \lambda u_{\langle s,e\rangle}.\lambda c.\lambda w.u(c)$ smokes in w

Ignoring tense and other such complications, the proposition that is expressed by an utterance of 'Mary smokes' in context c is derived using our composition rules as follows:

(11) $[\![\text{Mary smokes}]\!]^c = \lambda w.\text{Mary smokes in } w$

Suppose I don't know whether Mary smokes, but I learn from you that she does when I hear you say 'Mary smokes' in c. Then, according to the view I'm defending, I update my information state by eliminating those worlds that do not make $[\![\text{Mary smokes}]\!]^c$ true. For example, if my initial belief state $s = \{w_1, w_2, w_3, w_4\}$ and only w_1 and w_4 make $[\![\text{Mary smokes}]\!]^c$ true, then my posterior belief state will be $s' = s[\text{Mary smokes}]^c = s \cap [\![\text{Mary smokes}]\!]^c = \{w_1, w_4\}$. While the information gained is not context-sensitive, it is clear to see that my information state remains the same as I move across contexts, so long as I don't gain or loss any information. It continues to support the proposition expressed by the sentence 'Mary smokes' in c. And so the belief that I learnt from you is preserved.

Let us now see how this framework works for context-sensitive expressions by considering a context-sensitive fragment of English that includes pronouns. I will assume a referential analysis of the first person pronoun *I* according to which its intension is a function from contexts to the agent in that context, c_A (cf. Kaplan 1989):

(12) $[\![\text{I}]\!] = \lambda c.c_A$

Then, the Kaplanian character of (13) is as in (14):

(13) I smoke.

(14) $[\![\text{I smoke}]\!]^c = \lambda c.\lambda w.c_A$ smokes in w

Suppose Mary is the agent in c and John is the agent in c'. Then, the present analysis correctly predicts that the proposition that Mary expresses when she says 'I smoke' is different from the proposition that John expresses when he says the same sentence, since Mary refers to herself when she says 'I', and John refers to himself when he says 'I'.

We can now see the power of this proposal in explaining the preservation of content across contexts. For suppose now that w_1 is a world where both Mary and John smoke and w_4 is a world where Mary but not John smokes. Then my information state s' will be consistent with what Mary says in c, but not with what John says in c'. Once I learn that John smokes from what he said in c', I will update my information state accordingly and be left with $\{w_1\}$. Regardless of the context in which I find myself, what I learn from Mary and John's utterances is preserved in my information state, since the propositions that they expressed are both true in my singleton world. Of course, *I* cannot express my beliefs about Mary and John by saying the same sentence that they used, for the words 'I smoke' in my mouth expresses a totally different proposition entirely. But the inability to use their exact words to repeat what I learnt is no barrier for content preservation, once we see that its the information conveyed and not the form of the words used that were important to what is preserved.

5.4 Extending the account

Extending the framework further, let us consider gradable adjectives like *tall*. According to the contextualist account of gradable adjectives, the predicate *is tall* denotes the property of having a degree of height that is at least as great as the prevailing standard in the context. For concreteness, I assume that gradable adjectives have fixed denotations as suitable measure functions, and these combine with the null morpheme *pos* to form a Degree Phrase (DegP).[24] Thus, we have (15) for *tall*, where **tall** is a measure function that maps individuals to their height (at a world):

(15) $[\![[_{\text{Adj}}\ \text{tall}]]\!] = \lambda u_{\langle s,e\rangle}.\lambda c.\lambda w.\mathbf{tall}(u(c))(w)$

The meaning of *pos* is a function that takes gradable adjective denotations and returns intensionalised properties of individuals. So we have (16) for *pos*, where d_c is shorthand for the contextually appropriate standard comparison in c:

(16) $[\![[_{\text{Deg}}\ \text{pos}]]\!] = \lambda g_{\langle\langle s,e\rangle,\langle c,\langle w,n\rangle\rangle\rangle}.\lambda u_{\langle s,e\rangle}.\lambda c.\lambda w.g(u(c))(w) > d_c$

Thus:

(17) $[\![[_{\text{DegP}}\ [_{\text{Deg}}\ \text{pos}]\ \text{tall}]]\!] = \lambda u_{\langle s,e\rangle}.\lambda c.\lambda w.\mathbf{tall}(u(c))(w) > d_c$

[24] See, e.g., Bartsch and Vennemann (1972); Cresswell and Partee (1977); Kennedy (1999, 2007); von Stechow (1984).

Then, simplifying somewhat (and assuming that the copula is semantically vacuous), the Kaplanian character of (18) will be as follows:

(18) John is tall.

(19) $[\![\text{John is tall}]\!] = \lambda c.\lambda w.\textbf{tall}(\text{John})(w) > d_c$

To see how this analysis works, suppose Mary says the sentence 'John is tall' in context c, where the standard for tallness is $d_c = 170$. Then she expresses the proposition that John's height is greater than 170cm:

(20) $[\![\text{John is tall}]\!]^c = \lambda w.\textbf{tall}(\text{John})(w) > 170$

Suppose that John, who is ignorant about his height, is 175cm tall. According to the view under consideration, when John hears Mary say 'John is tall' in c, he learns something true about his height, namely, that his height is greater than 170cm. If he updates his information state accordingly, then for any world w' in John's updated belief state, his height in w' will be greater than 170 cm. That is, John's posterior information state will be $s' = s[\text{John is tall}]^c = s \cap \{w : \textbf{tall}(\text{John})(w) > 170\}$.

Suppose John later finds himself in a different context, c^*, where the standard for tallness is $d_{c^*} = 180$, somewhat higher than in c. According to the view, by moving to a new context, John loses the ability to truly report his belief state by using the sentence 'I am tall'. For if John tries to report on his height by uttering 'I am tall' in c^*, he will express a different proposition to what Mary expressed in c, namely, the proposition that his height is greater than 180cm:

(21) $[\![\text{John is tall}]\!]^{c^*} = \lambda w.\textbf{tall}(\text{John})(w) > 180$

The proposition that 'John is tall' expresses in c^* — the proposition that John's height is greater than 180cm — is different from the proposition that he learnt from Mary. Assuming that there are some worlds in John's information state in which his height is less than 180cm, John's previously updated information state will no longer accept 'John is tall' in c^*. However, while John has lost the ability to truly report his belief state with the sentence 'I am tall', he has not lost any information, since his information state hasn't changed from what he learnt in c. For the *proposition* that Mary expressed in c is true in both c and c^*, and his information state still accepts that proposition, that his height is greater than 170cm. Content is again preserved across contexts.[25]

[25] The argument that I am making in this section is fully general and can be extended to other context-sensitive expressions to explain the preservation of content. For some context-sensitive expressions and related phenomena, such as conditionals, propositional attitudes, and quantifiers, more sophisticated semantic machinery may be needed to handle other context-sensitive expressions. While it is outside the scope of this paper to provide those expression-specific resources, the general strategy can be easily adapted.

For other context-sensitive expressions, our representation of information states will need to be generalised. For example, epistemic modals arguably communicate facts about different agents'

5.5 Information, gained and lost

Testimony is not our only route to knowledge. Belief systems evolve in response to a variety of environmental pressures. Possibilities are eliminated by evidence gained through perception, memory, reasoning, and so on. Like all humans, John is continuously gathering and storing information about the context he is in, such as where he is, what time it is, what the weather is like, and myriad other contextual information. When John hears Mary say 'John is tall', he thus learns more than just the proposition she expresses. In particular, he learns that Mary has used certain words to communicate certain content and he stores that meta-linguistic information as well. By shifting from a view of belief as the preservation of *sentences* to a view of belief that preserves *semantic content* — or propositions — we do not lose the ability to store metalinguistic information altogether. That is, in addition to storing the proposition that his height is greater than the contextually salient standard, John will also store the proposition that Mary said the sentence 'John is tall' to communicate that content in a certain context. So if John finds that he cannot express in *his* context the same proposition that Mary expressed in *her* context, there are other means to express this old information, such as by reporting on the words that Mary used (e.g., 'She said 'John is tall"') or by expressing the old information anew (e.g., 'I am tall for a philosopher'). We should not be worried that we are not always able to use the same sentence from which we learnt a certain piece of information to express that information in a different context, since we have other resources to hand. Information gained.

Humans lack perfect recall. As our memories fade over time, the details of the past may escape us. But our cognitive limitations do not constitute an argument that natural language is not context-sensitive. Transformations on one's belief state are not one-way streets. As I forget that it was John who was speaking, my belief state widens to include the possibilities previously excluded, such as the possibility that it was Mary who was speaking. But while I may not remember *exactly* who was talking, I do not completely lose the information that I learnt. As previously ruled out worlds slip back into my belief state, I may lose my belief that John smokes, but I still know that it was either John or Mary who was talking, even if I don't know which. Just as the present view can model the acquisition of new information, through testimony as well as perception, as a narrowing of belief state, it can model the loss of old information, such as forgetfulness, as a widening of your belief state to include more worlds. Information lost.

---

information states, while deontic modals communicate facts about what is obligatory or permissible under certain realms of normativity. Then, to adequately model what happens when we accept information conveyed by sentence containing such expressions, we must keep track not only of our own information states, but also the information states of others as well as what is obligatory and permissible under different realms of normativity. One fruitful direction worth exploring is to generalisation the notion of an information states to a set of tuples, each constituent of which will consist in a set of worlds represents some parameter to be tracked. See Lewis (1979b), for some discussion of how this idea might be developed.

5.6 Context-sensitivity under uncertainty

Suppose I have amnesia and I can't remember whether I am John or Mary and someone is speaking to me, but I do not know who.[26] I am ignorant about what context I'm in. Luckily for me, the speaker makes the (rather bizarre) utterance 'I am either John or Mary, and I, but not you, smoke'. If I accept what they say, then I will have learnt (i) that either John or Mary is talking, (ii) that, whoever they are, they smoke, and (iii) that, whoever I am, I don't smoke.

Some help, but how much? My ignorance of who is speaking results in ignorance of what is expressed when the speaker says 'I smoke'. That is, I don't know whether it was John who expressed the proposition that John smokes or Mary who expressed the proposition that Mary smokes, although I will know that one of those two proposition were expressed. My ignorance of who is speaking results in ignorance of what is expressed when the speaker says 'You don't smoke': That is, I won't know whether I am Mary and I don't smoke, or that I am John and I don't smoke.

Some ignorance, but not much. I may be ignorant of what precisely was said, but I can often piece together enough information to narrow down which context I might be in. And from there, I can often piece together what propositions might have been expressed by the sentence, given what contexts I believe I could be in.

More precisely, to model uncertainty about the context, we need two things: a definition of a *cloud of contexts* and an update function that is sensitive to these clouds of contexts.[27]

Definition 5.4. A cloud of contexts C is a set of contexts. The set of clouds of contexts $\mathcal{C} = \wp(\wp(\mathcal{W}))$.

Once we have clouds of contexts, we can define a new update function, $[.]^C$, accordingly. Intuitively, atomic sentences narrow down an information state relative to a cloud of contexts by gathering together those worlds in the information state that make the sentence true relative to *at least one* context in the cloud, and eliminating the rest.

Definition 5.5. If $[.]^C$ is an update function at a cloud of contexts, C is a cloud of contexts, α is an atomic sentence, and ϕ, ψ are sentences, then:

$$s[\alpha]^C = s \cap \bigcup_{\forall c \in C} [\![\alpha]\!]^c$$
$$s[\neg\phi]^C = s - s[\phi]^C$$
$$s[\phi \wedge \psi]^C = s[\phi]^C[\psi]^C$$

To see how this new update function works, suppose that my initial information state is modelled as $s = \{w_1, w_2, w_3, w_4\}$, where both John and Mary smoke at w_1, John but not Mary smokes at w_2, Mary but not John smokes at w_3, and neither John

[26] Compare Stalnaker's remarks in his paper 'Assertion' (Stalnaker 1978).

[27] See von Fintel and Gillies (2008, 2011), for related discussion of clouds of contexts regarding epistemic modals.

nor Mary smoke at w_4. This information state usefully models my ignorance prior to the announcement. For all I know, I could be John and the speaker is Mary — call this context c_1 — or I could be Mary and the speaker is John — call this context c_2. Then, for all I know, the speaker could be uttering 'I, but not you, smoke' in c_1, and so express the proposition that John, but not Mary, smokes, or they could be uttering that sentence in c_2, and so express the proposition that Mary, but not John, smokes. Thus I update my information state relative to the cloud of contexts $\{c_1, c_2\}$ rather than just a single context, resulting in the posterior information state $s' = \{w_2, w_3\}$. While I may not know what proposition was actually expressed in a context — after all, I may not know exactly what context I am in — this is no hinderance to a contextualist semantics. Rather, on the view that I am advocating, this reflects the genuine consequences of the ignorance in which we frequently find ourselves. Moreover, on the present view, the contextual features of which we remain ignorant and their impact on our knowing what was said in context can be usefully modelled.

6 Concluding remarks

Context-sensitivity pervades natural language. But there is no pressing issue with preserving the information expressed by context-sensitive information at a context. We believe propositions, not sentences, and the content encoded in those propositions persist across contexts.[28]

References

Bartsch, R. and T. Vennemann. 1972. Relative adjectives and comparison. In *UCLA Papers in Syntax*, eds. P. Schachter and G. Bedell, Volume 2. Los Angeles, CA.: University of California Press.

Bezuidenhout, A. 2002. Truth-conditional pragmatics. *Noûs* 36: 105–134.

Borg, E. 2004. *Minimal Semantics.* Oxford: Oxford University Press.

Borg, E. 2012. *Pursuing Meaning.* Oxford: Oxford University Press.

Cappelen, H. and J. Dever. 2013. *The Inessential Indexical: On the Philosophical Insignificance of Perspective and the First Person.* Oxford: Oxford University Press.

Cappelen, H. and E. Lepore. 2005. *Insensitive Semantics.* Oxford: Blackwell Publishers.

Carston, R. 2002. *Thoughts and Utterance: The Pragmatics of Explicit Communication.* Oxford: Blackwell.

Cohen, S. 1999. Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives* 13: 57–89.

Cresswell, M. J. 1990. *Entities and Indices.* Number 41 in Studies in Linguistics and Philosophy. Dordrecht: Springer Netherlands.

Cresswell, M. J. and B. Partee. 1977. The semantics of degree. In *Montague Grammar*, 261–292. New York: Academic Press.

[28] I should like to thank Benjamin Brast-McKie, Matthew Hewson, and Lukas Lewerentz for helpful discussion on the ideas contained in this paper.

Davidson, D. 1973. Radical interpretation. *Dialectica 27*(3-4): 313–328.
Egan, A. and B. Weatherson. (Eds.) 2011. *Epistemic Modals*. Oxford: Oxford University Press.
Elbourne, P. 2008. The argument from binding. *Philosophical Perspectives* 22: 89–110.
Field, H. H. 1972. Tarski's Theory of Truth. *The Journal of Philosophy 69*(13): 347–375.
Field, H. H. 1978. Mental representation. *Erkenntnis 13*(1): 9–61.
von Fintel, K. and A. S. Gillies. 2008. CIA Leaks. *The Philosophical Review 117*(1): 77–98.
von Fintel, K. and A. S. Gillies. 2011. 'Might' made right. In *Epistemic Modality*, eds. A. Egan and B. Weatherson, 108–130. Oxford: Oxford University Press.
Fodor, J. 1975. *The Language of Thought*. New York, NY.: Thomas Y. Crowell.
Fodor, J. 2008. *LOT2: The Language of Thought Revisited*. Oxford: Oxford University Press.
Groenendijk, J. and M. Stokhof. 1991. Dynamic predicate logic. *Linguistics and Philosophy 14*(1): 39–100.
Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
Heim, I. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. PhD dissertation, University of Massachusetts, Amherst, MA.
Heim, I. 1983. On the projection problem for presuppositions. *West Coast Conference on Formal Linguistics* 2: 114–125.
Kaplan, D. 1989. Demonstratives. In *Themes from Kaplan*, eds. J. Almog, J. Perry, and H. Wettstein, 481–563. Oxford: Oxford University Press.
Kennedy, C. 1999. *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. Ph.D thesis, UCSC, Santa Cruz.
Kennedy, C. 2007. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy 30*(1): 45.
King, J. C. 2003. Tense, modality, and semantic values. *Philosophical Perspectives* 17: 195–245.
Kratzer, A. 2012. *Modals and Conditionals*. Oxford: Oxford University Press.
Kripke, S. 1972. Naming and necessity. In *Semantics of Natural Language*, eds. D. Davidson and G. Harman, 253–355. Dordrecht: Reidel.
Lewis, D. 1973. *Counterfactuals*. Cambridge, MA: Harvard University Press.
Lewis, D. 1979a. Attitudes de dicto and de se. *The Philosophical Review 88*(4): 513–543.
Lewis, D. 1979b. Scorekeeping in a language game. *Journal of Philosophical Logic 8*(1): 339–359.
Lewis, D. 1980. Index, context, and content. In *Philosophy and Grammar*, Synthese Library, 79–100. Springer, Dordrecht.
Lewis, D. 1996. Elusive knowledge. *Australasian Journal of Philosophy 74*(4): 549–567.
Magidor, O. 2015, December). The Myth of the De Se. *Philosophical Perspectives 29*(1): 249–283.
Montague, R. 1968. Pragmatics. In *Contemporary Philosophy. A Study*, ed. R. Klibansky, 102–122. Florence: La Nuova Italia Editrice.
Montague, R. 1970. English as a Formal Language. In *Linguaggi Nella Societa e Nella Tecnica*, ed. B. Visentini, 188–221. Edizioni di Communita.
Perry, J. 1977. Frege on demonstratives. *The Philosophical Review 86*(4): 474–497.
Perry, J. 1979. The problem of the essential indexical. *Noûs 13*(1): 3–21.
Schaffer, J. 2018. Confessions of a schmentencite: Towards an explicit semantics. *Inquiry*: 1–31.

Schiffer, S. 1992. Belief ascription. *The Journal of Philosophy 89*(10): 499–521.

Schiffer, S. 1994. A paradox of meaning. *Noûs 28*(3): 279–324.

Schiffer, S. 1995. Descriptions, indexicals, and belief reports: Some dilemmas (but not the ones you expect). *Mind 104*(413): 107–131.

Searle, J. 1980. The background of meaning. In *Speech Act Theory and Pragmatics*, 221–232. Dordrecht: D. Reibel Publishing Company.

Sperber, D. and D. Wilson. 1986. *Relevance: Communication and Cognition*. Oxford: Blackwell.

Stalnaker, R. 1968. A theory of conditionals. In *Studies in Logical Theory*, ed. N. Rescher, 98–112. Oxford: Oxford University Press.

Stalnaker, R. 1973. Presuppositions. *Journal of Philosophical Logic 2*(4): 447–457.

Stalnaker, R. 1978. Assertion. In *Pragmatics*, ed. P. Cole, 315–332. New York: Academic Press.

Stalnaker, R. 1984. *Inquiry*. Cambridge, MA: MIT Press.

Stalnaker, R. 2014. *Context*. Oxford: Oxford University Press.

Stanley, J. and Z. G. Szabó. 2000. On quantifier domain restriction. *Mind & Language 15*(2 and 3): 219–261.

von Stechow, A. 1984. Comparing semantic theories of comparison. *Journal of Semantics* 3: 1–77.

Travis, C. 1996. Meaning's role in truth. *Mind 105*(419): 451–466.

Veltman, F. 1985. *Logics for Conditionals*. Ph.D. thesis, University of Amsterdam, Amsterdam.

Veltman, F. 1996. Defaults in update semantics. *Journal of Philosophical Logic 25*(3): 221–261.

Williamson, T. 2005a. Contextualism, Subject-Sensitive Invariantism and knowledge of knowledge. *The Philosophical Quarterly 55*(219): 213–235.

Williamson, T. 2005b. Knowledge, context, and the agent's point of view. In *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, eds. G. Preyer and G. Peter, 91–114. Oxford University Press.

Williamson, T. 2020. *Suppose and Tell*. Oxford: Oxford University Press.