



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/87502/>

Version: Accepted Version

Article:

Santorio, P (2016) Nonfactual Know-How and the Boundaries of Semantics. The Philosophical Review, 125 (1). pp. 35-82. ISSN: 0031-8108

<https://doi.org/10.1215/00318108-3321721>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Nonfactual Know-How and the Boundaries of Semantics

Paolo Santorio
UNIVERSITY OF LEEDS

This paper is about three topics: know-how; metaethical expressivism; and foundational questions in the theory of meaning. Know-how and expressivism are usually regarded as disjoint subjects, belonging to distant provinces of philosophy. The investigation of know-how is taken to fall squarely within philosophy of mind, while expressivism has been invented and developed primarily as a position in metaethics. I argue that, despite obvious differences, these debates have structural similarities that run deep. In particular, one can make major moves on one side by mirroring moves made on the other. Semantic and conceptual tools developed by expressivists can be exported to the know-how debate and put to use in deflecting an influential line of argument about know-how. Moreover, expressivism provides the resources to create a new framework for thinking about know-how. This framework is nonfactualist—it validates the idea that knowing how to do something is different from knowing a fact—but overcomes problems that vitiate more classical nonfactualism, and especially the objection that nonfactualism about know-how is incompatible with our best semantics for know-how reports.

1 Overview

The contemporary debate on know-how starts with chapter 2 of Ryle's *Concept of Mind* (1949). There, Ryle argues that intellectual operations and skillful action are different kinds of activities and involve different kinds of mental states. Intellectual operations involve propositional mental states like beliefs and propositional knowledge;¹ skillful action requires know-how, which Ryle identifies with abilities. Since abilities don't have propositional content, know-how is not a kind of propositional knowledge. Ryle contrasts his position with the view he calls 'intellectualism', i.e. the claim that know-how

For comments and helpful discussion, thanks to Alex Byrne, Jennifer Carr, Yuri Cath, Dilip Ninan, Alejandro Pérez Carballo, Brian Rabern, Robert Stalnaker, Daniel Stoljar, many anonymous referees at *Philosophical Review*, and audiences at the ANU, the University of Leeds, the University of Rochester, the University of Illinois at Chicago, the University of Nottingham, and the Ordinary Language, Linguistics, and Philosophy Conference at the University of St Andrews (June 2011).

¹Throughout the paper I'll be assuming, following Williamson 2000, that there are factive mental states. Hence both propositional knowledge and know-how count as mental states. (On the nonfactualist picture, it's not immediately clear what the factivity of know-how amounts to. See section 4 for a partial answer.)

is a kind of knowledge-that. Here I don't discuss Ryle's argument, which much recent literature has shown to be problematic at best. Moreover, I don't rely on his characterizations of intellectualism and know-how, which are equally contentious.² My interest is in the basic categories introduced by Ryle and in the intuitive contrast between a kind of knowledge that represents facts and a kind of knowledge that guides subjects in action. Although many aspects of Ryle's original discussion are obsolete, the contrast between these two kinds of knowledge is still at the basis of the contemporary debate.

While Ryle's position was dominant throughout the twentieth century, the orthodoxy has been reversed over the past years. Stanley and Williamson (2001), and more recently Stanley alone (2011), have advocated the idea that, contrary to Ryle's account, someone's knowing how to ϕ just consists in their knowing certain propositions. Stanley gives a pithy and effective statement of the view:

[K]nowing how to do something is the same as knowing a fact. It follows that learning how to do something is learning a fact. For example, when you learned how to swim, what happened is that you learned some facts about swimming. Knowledge of these facts is what gave you knowledge of how to swim. (Stanley 2011, page 2).

In this paper, I'm going to use 'factualism' as a label for any view incorporating the claim that knowing how to do something consists in knowing facts, and 'nonfactualism' as a blanket term for views that deny it.

The argument that takes center stage in Stanley and Williamson's (henceforth, S&W) defense of factualism is based on the semantics of know-how ascriptions. They observe that know-how ascriptions like

(1) Sam knows how to cook risotto.

share both syntactic structure and compositional semantics with other knowledge reports involving embedded questions, such as

(2) Sam knows who cooked the risotto he ate.

It is an established claim in the semantics literature that sentences like (2) are just ascriptions of propositional knowledge. From here, they infer that, given the uniformity in the syntax and the semantics of the two sentences, also sentences like (1) work as ascriptions of propositional knowledge. Hence, they conclude, know-how is a kind of knowledge-that.

My basic observation is that S&W's semantic argument is interestingly similar to a classical objection to metaethical expressivism. Expressivism is a form of antirealism

²For a thorough discussion of Ryle's argument, see Stanley & Williamson 2001, Stanley 2011; for a critique of the identification of know-how and abilities, see Ginet 1975 and, more recently, Bengson & Moffett 2011.

about normativity, motivated (at least in part) by the rejection of normative facts.³ Contrary to other forms of antirealism, however, expressivists want to fully vindicate normative thought and talk. Roughly, they claim that normative statements don't express representational attitudes like beliefs and hence their contents are not ordinary propositions. One classical objection to this view is what goes under the name of 'Frege-Geach problem': expressivists are unable to account for the compositional structure and the logical properties of ethical claims. Thus there is a basic analogy between the two debates. In both cases, a metaphysical claim (know-how is not a relation between subjects of knowledge and facts, and there are no facts pertaining to the normative) is rejected on semantic grounds.

Sophisticated brands of expressivism develop a semantics that is designed to answer this worry. Here I take as my benchmark theory Allan Gibbard's semantics for norm-expressivism (1990). Throughout this paper, I explore what happens once we export to know-how the main insights behind Gibbard's moves in metaethics. I do this in two stages.

First, I argue that S&W's argument can be resisted via an expressivist strategy. Gibbard's central insight is that we can retain the basic compositional structure of standard semantics for normative language, while remaining neutral about what contents are expressed by normative claims. Similarly, I argue, we can fully help ourselves to standard compositional semantics for embedded questions, while remaining neutral about what mental contents are ascribed by know-how reports. The result is a treatment of know-how reports that mirrors, in relevant respects, the functioning of reports of normative attitudes on Gibbard's picture. In both cases, it is philosophy of mind rather than semantics that drives our views about the contents of states of know-how.

After blocking S&W's argument, I turn to the task of sketching a general view of know-how that is nonfactualist and at the same time is fully compatible with standard semantics for embedded questions. Again, the parallel with expressivism is fruitful. I define a new attitude, that of having a plan, which works as a non-factive counterpart of know-how. The semantics for having a plan is designed on the blueprint of possible world semantics, but exploits different kinds of atoms, namely *maximal performance plans*. Intuitively, these are fully specified sets of instructions for executing tasks. The difference in content between belief and having a plan is justified by a difference in functional role. The functional role of belief is (among other things) recording and storing information from the environment. I deny that having a plan has a functional role of this kind. Rather, its main functional role is guiding performance of action.

³This comes with a qualification. Proponents of expressivism (at least, of certain core versions of expressivism, such as Gibbard's norm-expressivism) insist that they can vindicate ordinary talk of facts. But their metaphysical picture is different from the realist's, at least in the sense that they are not willing to bestow on normative facts the same kind of metaphysical status that they grant to facts of (say) physics. It is this difference that I want to latch on to here.

The overall project of this paper connects to a trend of recent work in philosophy of language. Over the past ten years, a number of writers have defended brands of expressivism about epistemic, deontic, and probabilistic vocabulary that are explicitly inspired by Gibbard: see, among many, Swanson 2006 and 2011; Yalcin 2011, 2012a, 2012b; Moss 2014 (as well as her 2013 for some consequences about philosophy of mind); Rothschild 2012; Charlow 2013. This paper can be seen as sketching a related kind of expressivism about a different linguistic domain. The idea of using expressivistic tools for understanding know-how has not been explored before, to my knowledge. So most of the paper focuses on foundational work: motivating the view by showing how it connects to more traditional nonfactualism, and providing a general nonfactualist account of the contents of states of know-how. I sketch a formal semantics for know-how reports in the appendix.

2 The semantic argument for factualism

2.1 Stanley and Williamson's argument

Before starting, one *caveat*: I use Stanley and Williamson's original paper as my main source for the factualist line that I want to reject. I will refer explicitly to Stanley's more recent book whenever it makes a difference.

Let me begin by giving an explicit definition of factualism:

(Fact) S 's knowing how to ϕ consists in S 's standing in the knowledge-that relation to a certain relevant proposition or set of propositions.

Consider an example. After careful study and much experimentation in the kitchen, Sam has acquired the know-how that is relevant for making risotto. The proposition that Sam has learned is that *a certain way w is a way for him to make risotto*. Knowledge of this proposition is what guides Sam while cooking: for example, knowledge of this proposition is what makes it the case that he's able to tell when the rice has toasted long enough before he starts pouring in the broth.⁴

Notice that (Fact) is not meant to specify a full account of know-how. Factualists will likely want to supplement (Fact) with some more specific claims. For example, S&W claim that know-how involves entertaining propositions in a specific way or under a particular mode of presentation. Sam can get acquainted with the right way of making risotto by watching a TV show or by seeing me cook risotto for him. In these cases, he might come to believe the relevant proposition, but won't get the relevant know-how.

⁴If this sounds like an unfamiliar step in making risotto, check out the instructions at <http://culinaryarts.about.com/od/ricegrains/ss/risotto.htm>.

This, according to S&W, because the proposition is not entertained in the appropriate way. I won't be concerned with this aspect of their view for the moment, so I set it aside.

Let me introduce some terminology. I use 'knowledge-wh reports' to denote all knowledge reports where the complement clause is an embedded question. I also use 'know-how reports' to denote the knowledge reports with infinitival clausal complements that are typically employed in attributions of know-how. An example of a know-how report is

- (1) Sam knows how to cook risotto.

Now, to the argument. S&W start by providing a syntactic analysis of know-how reports. In line with standard views in linguistic syntax, they claim that sentences like (1) involve a complement clause with two covert features: a covert subject, the pronoun PRO, and a trace t that has the same subscript as *how*⁵:

Sam knows [how_i PRO to cook risotto t_i]

S&W's next move consists in noticing that this syntactic structure is shared by all knowledge-wh reports. Thus (1) shares its basic syntax with knowledge-wh reports with untensed complement clauses, like

- (3) Sam knows when to cook risotto.
(4) Sam knows where to eat good risotto in New York City.

as well as with knowledge-wh reports involving tensed complements, like

- (5) Sam knows who cooked the risotto he ate.
(6) Sam knows how I cooked risotto for him.

Finally, they point out that all these statements are treated on a par by semantic accounts of embedded questions. For example, on the account proposed by Karttunen 1977, all embedded questions denote the set of their true answers. Thus the embedded question 'who cooked the risotto he ate' in (5) denotes the set containing the true proposition x *cooked the risotto Sam ate*, where x is the person who actually cooked the risotto Sam ate. Similarly, the embedded question 'how to cook risotto' denotes the set containing the proposition w *is a way for Sam to cook risotto*, where w is indeed such a way.

On this basis, S&W conclude that know-how reports, on a par with other knowledge-wh reports, state that a subject has knowledge of a proposition, in the sense of 'knowl-

⁵The trace t just marks the position where the question word *how* was before undergoing syntactic movement. The subscripts, which are normally called *indices*, mark the fact that the two elements are linked syntactically and semantically.

edge' that is familiar from knowledge-that reports:

From a linguistic perspective, very little is special about ascriptions of knowledge how. It is hard to motivate singling them out for special treatment from the rest of a family of related constructions. Our view of ascriptions of knowledge-how is the analysis reached on full considerations of these constructions by theorists unencumbered by the relevant philosophical prejudices. (2001, page 431)

From here, switching from the formal mode to the material mode, they claim that know-how just consists in the obtaining of propositional knowledge.

2.2 Resisting the argument: two options

It's useful to lay down a schematic representation of the argument:

- (P1) *Semantic uniformity.* All knowledge-wh reports have a uniform compositional semantics.
- (P2) *Knowledge-wh as propositional.* Knowledge-who, knowledge-when, etc. reports are ascriptions of propositional knowledge.
- (P3) *Truth-conditional uniformity.* If all knowledge-wh reports have a uniform syntax and semantics, and if other knowledge-wh reports are ascriptions of propositional knowledge, then also know-how reports must be ascriptions of propositional knowledge.
- (C1) Know-how reports are ascriptions of propositional knowledge. (From (P1), (P2), (P3))
- (C2) Know-how consists in propositional knowledge. (From (C1))

The argument essentially consists of two steps. The first, from (P1)–(P3) to (C1), moves from the uniformity in the compositional semantics of knowledge-wh reports to the claim that their truth conditions are uniform. The second, from (C1) to (C2), moves from claims about attitude reports to claims about attitudes themselves.

Opponents of factualism generally resist the argument by rejecting the second step. The strategy is to resist the very idea that language should be relevant to the nature of mental states. Empirical work in cognitive science has vindicated a distinction that suggestively resembles the folk distinction between know-how and propositional knowledge (see, among many, Wallis 2008, Adams 2009, and references therein). In the face of this evidence, the argument goes, all arguments based on language are relics from the behaviorism-ridden era of ordinary language philosophy. Here is a typical statement of this view, due to Alva Noë:

Stanley and Williamson's investigation is in some ways methodologically backward. It is a mark of philosophical progress that we can now see that neither linguistic analysis nor cultivated intuitions are the key to understanding the nature of mind. (2005, page 290)

This is not the position I want to take here. In fact, I think that the step from (C1) to (C2) is unproblematic—or, at least, rejecting it commits us to an error theory that has high costs for philosophy of mind.

This claim is not central to my purposes, but let me briefly motivate it. I assume here that philosophical theories of attitudes aim at systematizing folk psychological notions with the final aim of interpreting and making sense of agents' behavior.⁶ Hence the notions of attitudes that we use in these theories are refinements of the corresponding folk psychological notions. Accordingly, ordinary ascriptions of attitudes are going to be part of the input data to our theory. Part of what we want to explain with a philosophical theory of belief, desire, or know-how is our ordinary ascriptions of beliefs, desires, and know-how.

If this is correct, it would be a significant cost if our theory of attitudes ended up systematically falsifying all ordinary ascriptions of know-how. A theory of know-how that treats our know-how ascriptions as systematically mistaken is a theory that fails to accommodate a big part of its starting data. For a comparison, take belief: a philosophical theory of belief that systematically falsified all belief ascriptions in natural language would be hard to take seriously. Admittedly, it is an empirical possibility that ordinary language ascriptions cannot ultimately be vindicated. But we can only embrace this possibility as a last resort. In the meantime, we should keep trying to square our philosophy of mind with ordinary thought and talk about know-how.⁷

Hence I take no issue with the step from (C1) to (C2). But I do think that S&W's argument is unsound. The problematic step is the one from (P1)–(P3) to (C1), and in particular premise (P3). As I'm going to argue, (P3) involves a conflation between different levels in a theory of meaning: structural analogies in the compositional semantics

⁶While this claim might not be shared by everyone, it is widely endorsed. For example, it is at the center of all versions of interpretivism, the view that the correct assignment of mental states to a subject is the one that allows us to make best sense of their behavior. (For a classical statement of this view, see Lewis 1974.)

⁷One worry, which seems to motivate Noë and others, is that factualism might seem incompatible with experimental evidence concerning know-how. By following S&W, the thought goes, we end up with a theory that is in conflict with our best cognitive science of know-how; this is clearly unacceptable. But the worry is misguided. There is no conflict between factualism and experimental results, at least if we stick to the definition of factualism that I've adopted in this paper. As Glick (2011) points out, questions about content—i.e., questions about the formal objects we use to index states of know-how—are independent from many questions typically investigated in empirical work. (I should point out that, while I agree with the main gist of Glick's paper, some of his claims about content might turn out to be in conflict with my position. It is crucial for me that some mental states be ascribed different kinds of content in virtue of those mental states having different functional roles. I'm not sure Glick would agree with this.)

don't guarantee analogies at the level of truth conditions.

3 How to Gibbard a Stanley-Williamson

3.1 Gibbard's maneuver

S&W's argument contains an important lesson. Any account of know-how should deal with the uniformity in the compositional semantics of embedded questions. I am in full agreement with them on this. But they fail to see that the compositional mechanics of know-how reports is not tied to any particular choice of denotation for embedded questions. On standard accounts, embedded questions denote sets of propositions. But we can modify this component of the account, thus assigning to questions semantic values of a different kind, while still using the basic structure of Karttunen-style semantics. This switch in semantic values paves the way to a nonfactualist account. A similar switch in semantic values is exemplified by expressivistic semantics; hence this is a good point to start developing the parallel with expressivism.

As I flagged in the introduction, a lot of recent work in philosophy of language has focused on developing expressivistic accounts of various kinds of discourse, in particular epistemically modalized discourse. My project is similar in spirit. But, rather than setting up a direct comparison with contemporary expressivistic views, I take as my starting point Gibbard's original norm-expressivism. I have two reasons. On the one hand, this makes it easier to see the main move I'm making. On the other, my focus here is on theoretical issues in philosophy of language that allow me to resist S&W's argument. Discussing recent expressivism would force me to introduce formal complications that are not central to my project.⁸

Gibbard's expressivism is a kind of *noncognitivism* about normative discourse. Gibbard holds that normative claims like 'Cannibalism is wrong' or 'Abortion is permissible' don't express beliefs, but rather conative attitudes of some sort. Following Gibbard 1990, I call these attitudes 'acceptances'.⁹ Accordingly, the contents of normative claims are not propositions: they are not truth-apt and are not meant to describe a way

⁸Yalcin (2012a) argues that expressivism is not a view about compositional semantics, but is better characterized as a view in the pragmatics. This is a good place to make it clear that I essentially agree with Yalcin's view—though, following MacFarlane 2005, I would prefer to use the label 'postsemantics' rather than 'pragmatics'. I do say that expressivism requires a change in the assignment of compositional semantic values, but this is compatible with Yalcin's view, as he himself makes clear. This change is necessary, but not sufficient to get expressivistic contents. This said, for ease of exposition I'll often use the label 'expressivistic semantics' to refer to the kind of semantic theories endorsed by expressivists.

⁹Gibbard has amended and updated his views in later work, most notably in Gibbard 2003. These updates are irrelevant for my purposes; what matters for me is the abstract way in which expressivistic semantics interacts with semantics for descriptive discourse.

the world is. Roughly, normative claims express endorsement or rejection of a normative standard. To claim that cannibalism is wrong is to express one's endorsement of a normative standard that prohibits cannibalism.

The decision to ascribe special contents to normative attitudes is driven by their different explanatory role in a theory of the mind. Beliefs are representational states: they purport to represent facts. Acceptances are conative states: they do not represent facts, but rather they contribute to fixing an agent's motivations and dispositions to act. Hence the ascription of special contents to normative attitudes is entirely determined by concerns in metaphysics (the absence of normative facts) and philosophy of mind (the motivational role of normative attitudes).

One of the challenges for the expressivist is showing how her view can yield a plausible semantics for normative language. Beliefs and acceptances are assigned contents of different kinds. But it would be disastrous if clauses like 'Cannibalism is wrong' and 'Cannibalism is widespread in New Jersey' were assigned different kinds of functions as semantic values. First, we would have to postulate systematic ambiguity in the expressions that interact compositionally with both clauses. For example, we should say that the two occurrences of *believe* in

- (7) Sam believes that cannibalism is wrong.
- (8) Sam believes that cannibalism is widespread in New Jersey.

have different meanings, i.e. denote two different functions.¹⁰ In (7), *believe* denotes a function taking propositions as arguments, in (8) a function taking normative contents as arguments. But even this wouldn't be enough. Descriptive and normative clauses can occur embedded together in a number of linguistic contexts. For example:

- (9) Sam believes that cannibalism is both widespread in New Jersey and wrong.
- (10) If cannibalism is both widespread in New Jersey and wrong, New Jersey people will be punished.

To accommodate (9) and (10), we should, first, decide on an assignment of a semantic value of the conjunction of a proposition and a normative content. Then, we should assume a further meaning for *believe*, different from the meanings used for (7) and (8), on which *believe* takes arguments of this kind. Even from these quick remarks, it should be clear that using different semantic values for descriptive and normative clauses leads to an intolerable multiplication of ambiguities.

This is the problem commonly referred to in the literature as 'Frege-Geach prob-

¹⁰Here I'm assuming that the expressivist will claim that natural language *believe* can be used for ascriptions of both beliefs in the stricter sense and Gibbard's acceptances. This move is obviously required to prevent expressivism from immediately turning into an error theory.

lem.’¹¹ Gibbard’s own answer to the challenge is simple. He adopts a general notion of semantic value that applies both to descriptive and normative clauses. He assumes as his benchmark theory possible worlds semantics, where all clauses denote sets of possible worlds (or, equivalently, functions from possible worlds to truth-values). The move is simply to generalize this notion of semantic value. On the new view, all clauses denote sets of pairs of a world and a system of norms to a truth-value. Systems of norms (for short, *norms*) are fully specified normative standards: they determine, for every possible act, whether it is forbidden, permitted, or mandated. In short, norms are a normative analog, at the formal level, of possible worlds. As a result of this move, descriptive and normative clauses are assigned semantic values of the same kind, i.e. sets of world-norm pairs.

$\{\langle w, n \rangle \mid \text{cannibalism is widespread in New Jersey in } \langle w, n \rangle\}$

$\{\langle w, n \rangle \mid \text{cannibalism is impermissible in } \langle w, n \rangle\}$

The conceptual insight behind Gibbard’s move is that metaphysical categories may be divorced from the categories we use in a compositional semantics. From a metaphysical standpoint, there are good reasons to draw a sharp divide between factual and normative information. But this distinction needs not be encoded in the semantics of natural language.

3.2 Rejecting uniformity

It’s easy to see how Gibbard’s maneuver can be brought to bear on S&W’s argument. Expressivist semantics teaches us that uniformity of semantic values in attitude ascriptions doesn’t entail uniformity in the kind of attitude contents ascribed. That is, in sentences like

(11) Sam believes that cannibalism is common in New Jersey.

(12) Sam believes that cannibalism is wrong.

¹¹Or at least, a part of the Frege-Geach problem. The literature on the topic is vast and complex, and often there is disagreement about what the problem itself is. For an overview, see Schroeder 2009, as well as Schroeder’s discussions in his 2008a and 2008b, among many others. I should note that it is controversial that Gibbard semantics offers a full solution to the Frege-Geach problem: Schroeder himself, in particular, has argued against this claim. I am inclined to think that it does and that Schroeder’s objections can be overcome, broadly for the reasons pointed out in Pérez Carballo 2012. If you’re less optimistic, please read my main claim as saying that the prospects of answering the S&W challenge to nonfactualism are connected to the prospects of solving the Frege-Geach problem for Gibbard-style expressivism. This claim, and the claim that there is a deep parallel between the two debates, are substantial and deserve attention even if you take the Frege-Geach problem to be an unsolved challenge for Gibbard.

the that-clauses have the same kind of semantic value. Nevertheless, their truth conditions involve ascriptions of attitudes of different kind, which play a different explanatory role in philosophy of mind. Similarly, one can grant, following standard semantics for questions, that the complement clauses in

- (1) Sam knows how to cook risotto.
- (3) Sam knows who cooked the risotto he ate.

have semantic values of the same kind. Yet there is still room for denying that they both state a relationship between a subject and a set of propositions. As the analogy with (11) and (12) shows, the type of the semantic values in play in attitude reports need not be a guide to the contents of the attitudes ascribed.

If this is correct, S&W's argument fails. Recall premise (P3):

- (P3) *Truth-conditional uniformity.* If all knowledge-wh reports have a uniform syntax and semantics, and if other knowledge-wh reports are ascriptions of propositional knowledge, then also know-how reports must be ascriptions of propositional knowledge.

Throughout this section, I have been arguing that the conditional in (P3) fails. Analogies at the compositional level in attitude reports do not guarantee analogies at the level of truth conditions. Even if we have a perfect compositional analogy, we must still look for validation in a general theory of attitudes. So (P3) is false and S&W's argument is unsound.

Let me dispatch a line of reply. One might protest that my reconstruction of the argument is not faithful to S&W's intentions. S&W don't start merely from assumptions about the semantics of questions (like my premises (P1)-(P3)). Rather, they assume the whole setup of existing semantic theories: this involves assuming that the basic semantic values of all clauses are propositions. On these assumptions, the objection goes, their conclusion does indeed follow.

I agree that S&W make this stronger assumption in their paper. I dispute that they're entitled to it. In a way, what I'm pointing out is precisely that they move illicitly from the weaker to the stronger assumption. Existing semantic theories were not designed to answer concerns about know-how. Assuming that we can read answers to questions in philosophy of mind straight off these theories as they happen to be currently set up would be unjustified.

3.3 Gibbard's maneuver, formalized

Before moving to my positive account, it's helpful to show how Gibbard's maneuver is implemented in a toy formal semantics. This will allow me to discuss and discard an im-

portant objection.¹² The discussion is accessible to readers with no formal background, and those uninterested in technical issues may skip ahead.

The basic task of a formal semantics is mapping each expression in a language to their denotations, usually called *extensions*, via an *interpretation function*, usually represented via the double brackets ‘ $\llbracket \cdot \rrbracket$ ’. Extensions are coarse-grained meanings: the extensions of referential terms are usually taken to be individuals, and the extensions of full clauses are truth values. For illustration, the following say that the extension of ‘Sam’ is the individual Sam, and the extension of ‘Sam is hungry’ is the truth value true.

$$\begin{aligned}\llbracket \text{Sam} \rrbracket &= \text{Sam} \\ \llbracket \text{Sam is hungry} \rrbracket &= \text{true}\end{aligned}$$

Usually, the interpretation function assigns extensions to expressions not absolutely, but rather relative to a series of parameters. For example, in many semantic frameworks, what extensions are assigned to expressions of the language depends on a world of evaluation.¹³ This relativization is represented via a ‘ w ’ superscript on the right-hand bracket. The following means that the denotation of the predicate ‘is hungry’, relative to a world of evaluation w , is the function mapping an individual to truth just in case that individual is hungry in w .

$$(13) \quad \llbracket \text{is hungry} \rrbracket^w = \lambda x. x \text{ is hungry in } w$$

(The lambda-notation is just a compact way to represent functions.¹⁴) The parameters to which interpretation is relativized in this way are usually collected in an n -tuple—the index of evaluation.

Index parameters have a double role in semantic theories.¹⁵ First, they are used to specify meanings that are more fine-grained than extensions. These are the meanings that, throughout this paper, I call ‘semantic values.’¹⁶ The semantic value of an expression is a function from index parameters to its extension. Using the straight brackets ‘ $\| \cdot \|$ ’ to represent a function from expressions to their semantic value, we have:

$$(14) \quad \|\alpha\| = \lambda i. \llbracket \alpha \rrbracket^i$$

¹²Thanks to an anonymous referee for raising this objection more than once and compelling me to take it seriously.

¹³Though this is by no means the standard treatment of modal parameters in contemporary frameworks. See Percus 2000 for a developed system which relocates reference to worlds in the object language.

¹⁴Here I’m assuming the definition provided by Heim & Kratzer 1998, section 2.5.

¹⁵For an in depth-discussion of the points I summarize in these paragraphs, see Lewis 1980.

¹⁶Semantic values are what is called ‘intension’ in some classical literature. Notice that, as Lewis 1980 pointed out, there are many plausible functions we might identify as semantic values. Functions from indices to denotations merely happens to be a popular choice.

Why define semantic values, if we already have extensions? Semantic values are crucial if compositional semantics has to provide an input to a theory of speech acts. Extensions are too coarse-grained to work as the objects that are asserted, presupposed, or believed. The extension of a sentence is merely its truth value. Hence, if extensions were what is asserted by uttering a sentence, there would be only two things we can ever assert—the true, and the false. By contrast, semantic values are, or allow us to recover¹⁷, more suitable objects—for example, possible worlds propositions.

The second role for index parameters connects to the lexical semantics of certain linguistic items: for example, modal auxiliaries like *might*. These items are usually modeled as ‘shifters’, i.e. as items that change the value of an index parameter. For example, *might* works by shifting the world of evaluation of the embedded clause to a different world.¹⁸

- (15) $\llbracket \text{It might be that Sam is hungry} \rrbracket^w = \text{true}$ iff there is an epistemically possible world w' such that $\llbracket \text{Sam is hungry} \rrbracket^{w'} = \text{true}$

The two roles of indices in semantic theories are easy to conflate, but (as we’ll see shortly) it’s important to keep them distinct.¹⁹

Now, in the terms of our toy framework, the expressivist’s maneuver consists simply in adding an extra parameter to the index of evaluation. For the case of normative discourse, following Gibbard, this is a norm parameter. Hence each expression is assigned a denotation relative to a pair of a world and a norm. Accordingly, semantic values turn out to be simply functions from a world and a norm to a truth value.

- (16) $\llbracket \text{Cannibalism is wrong} \rrbracket =$
 $\lambda \langle w, n \rangle. \llbracket \text{Cannibalism is wrong} \rrbracket^{w,n} =$
 $\lambda \langle w, n \rangle. \text{Cannibalism is wrong at } w \text{ and } n$

Different clauses will display different sensitivity to the world and norm parameter. The semantic values of normative clauses are world-insensitive: the world element in the pair doesn’t affect whether their semantic value maps world-norm pairs to truth or

¹⁷See footnote 19 for references about the relationship between semantic values and contents.

¹⁸This is obviously a toy semantics for modals. For an introduction to a developed formal semantics for modal operators, see von Stechow & Heim 2011, as well as the classical papers in Kratzer 2012. Just epistemic modals, as I remarked in the introduction, have become one of the main battlefields for expressivism. For relevant references, see section 1.

¹⁹ Arguably, the setup of standard theories is just based on this conflation. In *Demonstratives* (1989a), Kaplan explicitly identifies the so-called circumstances of evaluation (i.e. the parameters we use to define intensions) with the kind of parameters that are shifted by operators in natural language. Much literature has shown that this identification is a conceptual confusion, and should be resisted. For early statements of the point, see Dummett 1981 and especially Lewis 1980. For more modern defenses of the distinction see, among many, Ninan 2010 and Rabern 2012.

falsity. Conversely, the semantic values of factual clauses are norm-insensitive.

Now I can move to consider the objection. I state it as it applies to Gibbard's original view, but everything that I say holds, *mutatis mutandis*, for my semantics for know-how. The basic worry is that expressivistic semantics introduces extra complexity without appropriate justification. There are no compositional reasons to add a norm parameter to the index. But semantics for natural language should be exclusively driven by empirical concerns about the compositional assignments of meanings. Hence Gibbard's semantics, the argument goes, is gerrymandered and unjustifiedly complex.

The objection arises from a conflation of the two roles assigned to index parameters. One reason to introduce a parameter in the index is that we need it to model the semantics of certain lexical items. For example, one reason to introduce a world parameter in the index is that we need it for a compositional semantics for modal auxiliaries. Similarly, some philosophers have advocated a treatment of tenses that exploits a time parameter in the index.²⁰ If their empirical claims were right, we would have a reason to have a time parameter alongside a world parameter.

Crucially, though, this is not the only reason to add a parameter to the index. As I emphasized, index parameters have another role, i.e. allowing the theorist to define a suitably fine-grained notion of semantic value. What semantic values we need is not determined only by compositional interactions. Rather, it is determined by the kind of object we want to feed into a theory of speech acts.

To illustrate the point, I borrow an example from a recent defense of the very same point by John MacFarlane (2014, section 4.5). Suppose that you're giving a semantics for a primitive language that contains some basic vocabulary, but no modal operators. MacFarlane asks:

Would we take these speakers' lack of modal vocabulary to debar them from expressing the same kinds of propositions we express—for example, the proposition that snow is white? And would we say that, after they have acquired modal vocabulary, the contents of *all* of their beliefs change, and come to be true or false relative to worlds when they were not before? From the perspective of a philosopher of mind or theorist of speech acts, the idea should seem bizarre.

To strengthen MacFarlane's point: an anthropologist who learned the primitive language and acted as a translator between native and English speakers would, presumably, be preserving meaning at least to a very rough extent. But if our semantics for the primitive language can't help itself to an index parameter, no English sentence can be translated, not even approximately, into the primitive language.

²⁰The question whether interpretation should be relativized to times has been the subject of a long debate. The position that we should have a time parameter was initially adopted by Kaplan (1989a, 1989b); for criticism of that position, see, among many, King 2003, Kusumoto 2005.

The reason why we want to relativize interpretation to worlds in MacFarlane’s example is that our best theory of assertion uses contents that distinguish between ways the world might be. This constraint is not motivated compositionally, but it is still driven by empirical considerations. We want to give the best interpretation of utterances in the primitive language, and this requires using contents that cut across modal space. Gibbard’s argument for more fine-grained contents is analogous. For the expressivist, there are no normative facts. Yet we want to account for the fact that normative utterances play an important role in our cognitive economy. The solution is to assume that normative utterances make distinctions not between ways the world might be, but between normative possibilities. To derive these contents from the semantics, we assume that interpretation of sentences is relativized to norms besides worlds, and that hence semantic values are more fine-grained than on the standard picture.

Hence there is a straightforward argument for introducing norms in the index that is independent of compositional considerations. Before moving on, let me flag that I don’t think that this is the only argument we will find. Presumably, a fully developed expressivistic semantics for normative language will *also* assign a compositional role to the norm parameter—for example, to handle the semantics of deontic modals.²¹ Similarly, on the developed version of my semantics for know-how (in the appendix), some elements of infinitival questions selectively manipulate an extra parameter. But it’s important to emphasize that (if the expressivist is right in her claims about the nature of normative attitudes) we already have decisive reason to set up the semantics differently.

4 Know-how as a directive attitude

Let me take stock. S&W point out an important fact: the compositional semantics of know-how reports is fully parallel to the semantics of other knowledge-wh ascriptions. This must be accounted for by any plausible theory of know-how. But this doesn’t entail the truth of factualism. As expressivist semantics shows, we might acknowledge the uniformity in the compositional semantics while denying the uniformity of contents ascribed by knowledge-wh reports.

Even though their argument is unsound, it might still be that the conclusion is correct. Moreover, the burden of proof seems to be on the nonfactualist. S&W exploit the standard setup of semantics for embedded questions. It is up to the nonfactualist to propose a new notion of semantic value that can both fit into the standard compositional machinery (so that the uniformity premise invoked by S&W is vindicated) and allow us to treat know-how differently from propositional knowledge. From now on, I take up the task of developing a kind of nonfactualism that measures up to this demand.

²¹See Yalcin 2012a for an attempt in this direction.

Two qualifications are in order. The first is that I won't be offering a general defense of nonfactualism. Of course, if my account is viable, it shows that nonfactualists are able to vindicate our thought and talk about know-how. This can be seen as an indirect argument for nonfactualism. But a proper defense of nonfactualism is best left to another occasion.

Second, I won't be giving a philosophical analysis of know-how. My main purpose is outlining a credible alternative to factualism and explaining how a view of this sort should model the content of know-how. This doesn't amount to giving an analysis. In fact, the theory exploits a primitive notion, that of following an instruction, which seems no more intuitive than the notion of know-how itself.

4.1 A directive state

Here is the basic suggestion. Knowing how to ϕ consists in being in a mental state that reliably guides one to successful completion of a task. The content of this mental state is directive: it can be modeled as an instruction, or a set of instructions, for completion of the task. For example, Sam's knowing how to cook risotto consists in Sam's being in a mental state that reliably guides him while he's cooking risotto. The content of this mental state is a set of instructions detailing the operations that Sam performs to cook risotto. This whole section is devoting to fleshing out this idea and making it precise.²²

Let me clarify what I mean by saying that know-how has directive content. Contents mark the role of an attitude in a general picture of the mind. Once again, the analogy with expressivism is illuminating. The switch to non-descriptive contents is meant to capture two features that, according to the expressivist, normative attitudes possess. First, for the expressivist normative attitudes are nonrepresentational: their contents don't depict facts and don't have truth conditions in the standard sense. Second, the expressivist takes normative attitudes to have a special motivational force, which distinguishes them from beliefs. Hence the switch in content marks a switch in the functional role assigned to normative attitudes. Similarly, the claim that know-how has directive content means that know-how has a different functional role from propositional knowledge and belief. What is this new functional role?

²²The idea that know-how has a kind of directive content is very natural. Unsurprisingly, it has resurfaced again and again in the literature on know-how. Ryle 1949 himself entertains the suggestion (then rejected as a part of the "intellectualist legend") that know-how might involve "prescriptions". A notion of "procedural" or rule-based knowledge which was meant to cash out the traditional notion of know-how was formulated in the artificial intelligence literature (see, among many, Winograd 1975 and Cohen & Squire 1980).²³ David Carr has linked know-how to practical rationality, arguing that possessing know-how amounts to being instructed "by means of practical directives" (Carr 1981). My proposal is not a descendant of any of these accounts, but the popularity of the directive idea suggests that it embodies an important intuition.

Start from a rather crude picture: think of individuals as complex functional systems with upstream and downstream connections to the environment. Information enters via the upstream links; the downstream links result in behavior. I assume that propositional attitudes like beliefs and desires are partly constituted by the functional role they play in the processes that start with uptake of information and end in production of behavior. For example, part of what it is to be a belief is to stand in the right kind of functional connections to the individual's upstream and downstream connections to the environment, as well as to other kinds of mental states.²⁴ In particular, the functional role of representational states like belief will be (among other things) recording and storing information from the environment. Since recording and storing information involve a causal component, this means that representational states have—in the terminology of Stalnaker 1984—a backward-looking aspect. They are partly constituted by their standing, in normal conditions, in appropriate causal connections with the environment.²⁵ Notice that this is not to say that having a causal connection to the environment in normal circumstances is the only functional role of beliefs. On the contrary, philosophers agree that beliefs also have a forward-looking functional role. But my claim, which I take to be uncontroversial, is that the backward-looking aspect is necessary for a mental state to count as a belief.

What I deny is precisely that know-how has a similar backward-looking functional role. Representing the environment is not part of what it is to be a state of know-how. States of know-how don't have the function of standing in backwards causal connections to states of the environment. Rather, the central functional role of know-how is determining behavior, and in particular guiding performance of particular tasks. Know-how has a central forward-looking functional role, but not a backward-looking one—at least, not the same that is in play for belief.

This idea has two interesting consequences. The first is that it vindicates the link between know-how and action guidance. It is a truism that know-how is what guides expert performance of a task. This is, intuitively, one of the features that sets know-how apart from propositional knowledge. Possession of the relevant propositional knowledge is not sufficient for skillful performance. I might have detailed knowledge of the way people ride bikes, yet still be incapable of riding a bike in a wide variety of actual and counterfactual circumstances. Hence my having propositional knowledge is, by itself, not enough for me to have know-how.²⁶ Factualists must find a way to bridge this

²⁴Notice that this is much weaker than assuming that attitudes like belief and desire can be *analyzed* in functional terms. All I require is that, in order to count as a belief, a mental state must, possibly among many other things, have functional connections of a certain kind.

²⁵The 'normal conditions' clause is required to take care of misrepresentation cases and is notoriously hard to unpack. Incidentally, notice that causation might not be quite the right factor: other accounts (e.g., Dretske's 1981), use nomological covariation instead. This difference is irrelevant for my account.

²⁶Notice that the argument does not rely on the controversial identification of know-how with abilities.

gap. As I mentioned, Stanley and Williamson appeal to what they call ‘practical modes of presentation’: a subject knows how to ϕ just in case she knows the relevant proposition under a practical mode of presentation.²⁷ By contrast, I have no need for this extra bridging element. On my picture, possessing know-how just is to be in a state that is action-guiding. The connection between know-how and action guidance is built into the definition of the former. This, of course, doesn’t mean that the account explains the action-guiding role of know-how. But it does manage to avoid an appeal to modes of presentation or similar bridging notions.

The second consequence is that a forward-looking account of know-how seems to allow for lucky know-how. It is often observed that know-how seems immune to Gettier phenomena. To make the point, I borrow an example from Cath 2012²⁸:

Charlie wants to learn how to change a light bulb. He consults *The Idiot’s Guide to Everyday Jobs*. There he finds accurate instructions to perform the task and grasps them perfectly. As a result, he comes to be in a position to reliably change lightbulbs (in normal circumstances, and *ceteris paribus*). But, unbeknownst to Charlie, the guidebook has been written with the intention to misinform the reader and contains extremely misleading instructions for all other tasks. The accurate instructions are there because of a fluky computer error that caused random text to appear on just one page of Charlie’s copy of the book.

Despite the extreme flukiness involved in the case, it seems obvious that Charlie comes to know how to change a lightbulb as a result of reading the book. This brings out a major difference between know-how and propositional knowledge: obviously, Charlie is in no position to gain propositional knowledge from the book.²⁹

All that I need is that possession of know-how *entails* possession of ability in a suitable range of possible circumstances (not necessarily close-by ones). This much should be uncontroversial.

²⁷On the view defended in his 2011, Stanley drops the commitment to modes of presentation in favor of the idea that know-how reports exploit a special flavor of modality, which is slightly different from that of the overt modal *can*. (For the claim that know-how reports in natural language involve modality, see the appendix.) Thus, when Sam watches a TV show about making risotto without thereby acquiring the related know-how, what he learns is just a *different proposition* from the one whose knowledge would count as possessing know-how. This move might be an improvement on the S&W proposal in that it dispenses with the need for modes of presentation. But it still leaves the most pressing questions unsolved: what generates the action-guiding properties of these special propositions? And why can’t we express these propositions in speech, so that we can impart know-how simply by uttering a sentence, or grasp these propositions simply by watching TV?

²⁸For a similar case outside the know-how literature, see Pettit 2002.

²⁹*Caveat*: the claim that know-how is Gettier-immune is disputed by factualists. S&W (2001, page 435) present an example structurally analogous to the Charlie one, claiming that it is merely a case of justified true belief. But their example is unconvincing. Subjects that have no stakes in the know-how debate, including philosophers, overwhelmingly converge on the idea that all Charlie-like cases genuinely involve

A full explanation of Gettier immunity goes beyond the purposes of this paper. But it should be clear that a forward-looking account of know-how is well placed to give this explanation. As a backward-looking state, propositional knowledge requires a reliable connection of some sort between the subject and the world. You cannot know that the lightbulb is broken—you cannot have a backward-looking state in good standing—without a suitable connection to the fact that the lightbulb is broken. Gettier cases illustrate just the failure of this connection. But forward-looking mental states are exempt from this requirement. They are not in the business of representing facts in the world and their backward causal connections to the environment are irrelevant to whether they do their job properly. So there can be forward-looking states in good standing that are brought about by accident, as the Charlie case shows.³⁰

4.2 Performance plans

I now proceed to my main task, namely specifying a nonfactualist theory of content for know-how. In the next paragraphs, I give a full semantics for mental states. I will give a brief sketch of a semantics for know-how reports in language in the next section. A fully developed semantics, which is more technically involved, is postponed to the appendix.³¹

I start by defining an attitude that has the same functional role of know-how but differs in that it doesn't require success in performance. I call this attitude *having a performance plan*. Having a performance plan is essentially a nonfactive analog of know-how. I won't try to provide an analysis of this notion (I am indeed skeptical that analyses of this kind can be provided). Rather, I gloss it in terms of another notion, the notion of *an agent's behavior being governed by instructions*:

A subject *S* has a performance plan to ϕ in way *W* just in case (*ceteris paribus*) *S*'s behavior is governed by instructions according to which acting in way *W* is conducive to ϕ -ing.

know-how. In reply, Stanley (2011, chapter 8) argues that our intuitions about these cases are unreliable and that there is data suggesting that all knowledge-wh ascriptions are heard as Gettier immune for pragmatic reasons. This seems a strong claim and I don't think that Stanley gives a large enough survey of the data to support it. In any case, if intuitions in Charlie-like cases have pragmatic roots, as Stanley claims, surely it should be possible to find an example where context doesn't mislead us. But, to my knowledge, no one has produced a convincing example of this kind.

³⁰Incidentally, this point suggests that we might find Gettier analogs of know-how by looking at the *downstream* connections of a state with the environment. For reasons of space, I must set aside the question whether there are such cases.

³¹The model I propose is rather loosely inspired by Gibbard's (2003) semantics for deliberative discourse, namely discourse about what to do. I won't be going through Gibbard's original semantics, but the analogy will be transparent to readers familiar with his work.

The *ceteris paribus* qualification is there to screen off cases where something is temporarily awry in the agent's cognitive architecture (they're intoxicated, or extremely fatigued, etc.), or cases where they lack the ability to follow the instructions because of contingent reasons (say, their hands are tied).

A word of warning about terminology: the notion of having a performance plan only bears a limited resemblance to the ordinary notion of having a plan. Having a performance plan is not, in general, a personal-level attitude: all competent bike riders have a performance plan to turn the handlebar to the left at the beginning of a right turn, yet few of them are aware of this plan. By contrast, it seems strained to say that an agent has a plan, in the ordinary sense, to (say) spend their holidays in Croatia, yet they are unaware of it. Also, agents may have performance plans that are in no way manifested in bodily behavior—for example, agents might have a performance plan to engage in certain mental operations to solve a math problem. Despite these shortcomings, the label 'plan' still seems to me the best that I can find.³² But the reader should bear in mind that performance plans are not plans in an ordinary sense.

Let me say more about the notion of an agent's behavior being "governed by instructions". Consider an agent exemplifying some ordinary instances of know-how—say, Sam, who is a construction worker with a penchant for tennis and Italian cuisine. I say that, when Sam drives to work, lays bricks, hits a ball with a backhand stroke, or adjusts the fire under the risotto, he engages in behavior that is governed by instructions. This involves assuming that the explanation of Sam's behavior in all these circumstances will appeal to psychological states of a certain kind. These psychological states are part of the broadly functional picture outlined in the previous pages: when the appropriate functional connections hold, they will result in determinate patterns of behavior. Now, at any given time, we can think of the totality of the states of this kind that Sam instantiates. Hence, at any given time, we can talk about the overall set of instructions that guide Sam's behavior at that time. Right now this set includes (say) instructions for driving, laying bricks, hitting a tennis ball, etc. The set might expand or contract if Sam learns or forgets how to perform tasks.

The notion of an agent's behavior being governed by instructions is the keystone of the theory and it is the one that I take as primitive. It is obviously a philosophically loaded notion. So I do have a substantial primitive in my framework. This is not a worry: my goal is providing a framework to model the content of states of know-how, rather than trying to give a general theory of know-how.

Let me add three clarifications. First, the notion of agent's behavior being governed

³²An anonymous referee suggests 'strategy' or 'technique' as alternative labels to 'plan'. I agree that these lack some connotations that I want to leave out from the notion of a performance plan; at the same time, though, they also seem to lack the suggestion that having a performance plan is action-guiding, which is an important part of my proposal. So I have decided to stick to my label.

by instructions is a dispositional and not an occurrent one. What instructions govern a subject's behavior is a matter of what states would guide the subject during the performance of a task, should they engage in those tasks. Relatedly, having a performance plan is independent of having an intention to carry out the relevant instructions. Having a performance plan is being in a state such that, when it enters the right kind of functional connections, it guides performance of tasks. But that state might just never enter the relevant functional connections. A subject having a plan but lacking an intention to execute the plan exemplifies just this situation.

Second, I am construing instructions as linked to particular goals. In other words, instructions are never categorical (' ϕ !'), but always conditional on a certain goal (' ϕ , in order to satisfy goal G !'). For example: Sam's behavior is governed by instructions to act in such a such a way *in order to make risotto*; to perform such and such bodily movements *in order to hit a ball with a backhand stroke*; to handle bricks in such and such a way *in order to build a wall*; and so on. Also, I allow that a subject's behavior may be governed by multiple sets of instructions for performing a task and hence that a subject may have multiple plans to perform a task. (Though, of course, normally at most one set of instructions at a time will be in execution.) This is as it should be. The same subject may have the ability to ϕ in different ways. Mastery of one way doesn't exclude mastery of the others.

Third, I attach no significance to the glosses of the relevant instructions being in the imperative mood.³³ I claimed that having a plan and know-how have directive content. It seems natural to infer that the contents of these mental states are, and should be reported as, imperatives. But, as I emphasized in section 4.1, the claim that these states have directive content is merely a claim about their functional role. Roughly, having a performance plan has the functional role of guiding behavior when the subject engages in specific tasks. I want to remain neutral on the issue whether the content of know-how is literally some kind of imperative. The reason is that it's unclear to me how, if at all, the distinction between indicative and imperative mood may be transposed to the level of attitude content. (It is also controversial what the distinction amounts to in the first place; see Charlow 2014 for an overview.) This is an interesting question, but one about which I make no commitments at this stage.

Let me also address two objections. First, one might worry that the notion of having a performance plan is too demanding. The way I explicate having a plan seems to require thinking that one's plans will generally be successful, since the relevant instructions have to represent the course of action as conducive to bringing about the plan. But agents seem to have performance plans even when they are not confident in their own success. For example, a professional baseball player seems to have a performance

³³This is why, in the in the definition above, it is okay to gloss the content of the relevant instructions in the indicative mood (as "instructions according to which acting in way W is conducive to ϕ -ing).

plan to act in way W to bat, even though the success rate for batting among professional players is lower than 50%.³⁴ The reply is that what is represented in performance plans may come apart from what is represented in the agent's cognitive attitudes. Hence the MPPs compatible with a baseball player's plans may represent moving in way W as conducive to batting, while at the same time their credence distribution assigns less than .5 credence to worlds where they bat successfully. This is not a problem. Performance plans and credences are different kinds of psychological states. We should not expect their content to line up, aside from special cases.

Conversely, one might worry that my view makes success too easy. Knowing how to ϕ (and hence having a plan to ϕ in a certain way) doesn't entail being able to ϕ .³⁵ But, if an agent's behavior is guided by appropriate instructions to perform the task, then it seems that the agent will necessarily perform the task successfully. So, the worry goes, my view leaves no gap between knowing how and ability.

The objection relies on the following assumption: if an agent's behavior is governed by instructions that tell them to act in a certain way, they will necessarily act in that way. But I deny this. An agent's behavior might be governed by certain instructions and, at the same time, various factors might get in the way of the instructions being carried out successfully. Their peripheral neural circuitry might misfire. Their muscles might be too worn out or too weak. Their body might not comply in some other way. Take a concrete example: Sam has a performance plan to act in way W to hit a lob shot in tennis. This plan is a good one—the kind of plan that, in many circumstances, leads to success. As a result, as long as he's young and his body responds properly, Sam manages to hit successful lob shots pretty consistently. But now, suppose that, as Sam ages, his muscles become less responsive and he loses the ability of hitting lobs. All his attempts fail. Nevertheless, while he tries to hit lobs, his behavior is still governed by the same psychological states. His aging has had effects on his muscles, but not on his brain. Hence he still has the same performance plan and, presumably, he still knows how to hit a lob. He's just not able to do so any more.³⁶

This objection is rather vexing for other brands of nonfactualism (for example, the account in Hawley 2003). But my account handles it easily. Let me stress why: like the factualist, I am assuming that states of know-how are states individuated in part by a certain content, rather than just by behavioral dispositions. This allows me to recognize circumstances where the relevant mental states are present, while the behavioral outputs

³⁴Thanks to an anonymous referee for raising this issue and suggesting this example.

³⁵For a defense of this claim, see Bengson et al. 2009 and Bengson & Moffett 2011.

³⁶In the limit case, Sam might acquire an 'appropriate' performance plan to ϕ without ever developing the ability to ϕ . An alleged case of this kind is presented in Bengson et al. 2009: they describe a ski instructor who is able to successfully teach a number of skiers how to perform certain complex stunts, but unable to perform the stunts himself. When presented with this case, subjects judge that the instructor does know how to perform the stunts. I agree that this intuition is available and my account can vindicate it.

are not. This is one of the explanatory advantages of factualism that I can replicate, while still remaining on the nonfactualist side of the divide.

4.3 Maximal plans semantics

I now turn to the development of the formal framework. The main move is to treat the semantics of having a plan according to the blueprint provided by possible worlds semantics for belief. In the latter, we define a space of maximal ways the world might be, or possible worlds, and model the content of individual attitudes as sets of worlds, namely the worlds compatible with the attitude in question. Similarly, I introduce ‘total plans’, and I model the content of having a plan as a set of total plans. More precisely, I define a space of *maximal performance plans*. In informal terms, a maximal performance plan (henceforth, *MPP*) is a mapping of possible courses of action—by which I mean, simply: possible sequences of actions performed by the subject—to goals. This mapping will, in general, be many-to-one. A maximal performance plan is *total* in the sense that it maps every possible course of action to a goal. In metaphorical terms, you can think of an *MPP* as a gigantic lookup table linking courses of action with goals.³⁷

MPPs lend themselves to a formal definition. We can model a course of action simply as a set of worlds, i.e. the sets of worlds where that course of action takes place.³⁸ Similarly, we can model goals as the sets of worlds where the goals are achieved. Hence an *MPP* can be modeled as a (many-to-one) mapping of sets of worlds to sets of worlds. The mapping will capture a kind of causal relation: *MPPs* model what goals are brought about by certain courses of action.³⁹

So much for the basic elements of the theory. At this point, I have the resources to give a specification of content in this framework. In analogy to what happens with possible worlds and beliefs, the content of a plan is specified in terms of compatibility with the subject’s overall plans. More specifically:

³⁷Let me emphasize that, by introducing *MPPs*, we are not committing ourselves to unrealistic assumptions about what is represented in subjects’ minds. *MPPs* are formal tools that the theorist uses to model subjects’ mental states and are not supposed to be psychologically real. The same point applies to worlds in possible worlds semantics: see (among many others) the postscript to Lewis 1979.

³⁸For well-known reasons relating to self-locating attitudes (cf. Lewis 1979, Perry 1979), we might have to use centered worlds rather than possible worlds here. This point is not central to my purposes, so I set it aside.

³⁹It might seem implausible that every course of action is mapped to a unique goal. For example, one might lower the heat under the risotto with two goals: avoid burning the rice and cooking the vegetables at the right temperature. We can take care of cases of this sort by construing the goals appearing in *MPPs* as maximally specific, i.e. by taking conjunctions of what we ordinarily consider to be different goals. One might still worry that different courses of action can be undertaken with different and perhaps incompatible goals. I think this problem can be avoided by individuating courses of action in a more fine-grained way. If this was not enough to solve the worry, I’d be happy to construe *MPPs* as many-to-many mappings. Nothing in what I say requires them to be many-to-one.

S has a plan to ϕ in way W iff, for every MPP m compatible with S's plans, according to m acting in way W brings about ϕ -ing.

For illustration, take the usual risotto example. I say that Sam has a plan to cook risotto in way W —which, say, consists in throwing rice in a pot and slowly adding water until the rice is cooked—just in case, according to all MPPs compatible with Sam's plans, acting in way W is conducive to making risotto. Notice that the relevant MPPs might also allow other courses of actions that bring about the making of risotto; moreover, they need not agree on whether these other courses of action are conducive to risotto making. What matters is that they all agree that the course of action individuated by W invariably brings about Sam's cooking risotto.

Two clarifications are in order. First, I take the talk of acting in certain ways to be just a notational variant of the talk of courses of action. I appeal to ways just because they turn out to be useful for a semantics of know-how reports. Second, I should explain what it is for an MPP to be compatible with a subject's plan. As in possible worlds semantics for belief, this notion of compatibility is the primitive notion of my formal model. But, given what I've said so far, the basic idea should be pretty clear. An MPP represents what courses of action are conducive to what goals. Hence an MPP is compatible with a subject's plans just in case those plans don't rule out a certain course of action being conducive to a goal. For example, an MPP that maps flapping one's arms to the goal of riding a bike is compatible with a subject's plans just in case the subject's overall plans don't rule out flapping one's arm being conducive to successfully riding a bike.

Possible worlds semantics for belief models not only static beliefs, but also learning. When a subject learns a proposition p , the set of worlds that are compatible with their belief state is shrunk by ruling out all worlds which don't validate p . This feature of the framework carries over to plans and MPPs. Acquiring a plan can be modeled by shrinking the set of MPPs that are compatible with the subject's plans. Suppose that Sam undergoes training for cooking risotto. If Sam is completely ignorant at the beginning of his training, all sorts of MPPs will be compatible with what he knows. For any course of action A , some MPPs will represent A as conducive to making risotto, others won't. As Sam starts learning, some of these MPPs are ruled out. At the end of the training, all of the MPPs compatible with Sam's plan will agree in representing some course of action as conducive to risotto-making. If the training is indeed successful, this will be in fact a good course of action for making risotto; but of course, it might be that this course of action actually leads to failure and that his plans don't amount to possession of know-how.

4.4 From plans to know-how

Having a plan is meant to be a non-factive analog of know-how. It seems clear that there should be such an analog. There are cases where a subject performs a task on the basis of action-guiding states, whether successfully or not, yet she doesn't count as having the relevant know-how. At the same time, it's not immediately clear what the factivity of know-how amounts to.⁴⁰ It's obviously too strong to say that know-how requires a subject to *always* perform the task successfully, even when all enabling conditions are in place (the subject is not intoxicated, or too tired, etc.). Here is a more promising idea: we require that, to possess know-how, a subject be able to reliably perform the task in a suitable range of circumstances. I take this line here. So I endorse the following analysis of knowing how to ϕ :⁴¹

S knows how to ϕ in way W in context c iff, in c , S has a plan to ϕ in way W and this plan is reliably successful across circumstances C .⁴²

Notice that this analysis exploits a number of parameters: a reliability threshold, a notion of success, a range of possible circumstances. Interestingly, it seems that the values of these parameters need not be set once and for all in all cases. We can get know-how in circumstances where the values of these parameters are set in very different ways. Let me give some examples.

First, there might be variation in the circumstances that are taken as suitable for evaluating possession of know-how. Suppose that, after some training, Sam has acquired the ability to make risotto by using his own kitchen tools and a certain kind of ingredients. But he would miserably fail in the task if he were to use different tools or different ingredients; his expertise just isn't broad enough yet. In this kind of circumstance, it seems that we might ascribe Sam knowledge of how to make risotto, but also we might not. Second, there is variation in what counts as successful performance of a plan. If Sam reliably cooks a barely edible risotto, we might or might not count him as having the relevant know-how. Third, we might have variation in how reliably successful one has to be to count as having know-how. It seems plausible both that Babe Ruth

⁴⁰For another attempt (in a different vein, and with different aims) at extending the standard notion of factivity, see Moss 2013.

⁴¹There are clear similarities between this analysis and the account defended by Hawley 2003, though I explicitly deny that the suitable circumstances need be a set of closest worlds. The key difference between Hawley's account and mine is that I construe know-how as involving a psychological state with content. She, by contrast, just uses a notion of counterfactual success in action, combined with an epistemological notion of warrant. Just the appeal to content, I believe, is what allows me to escape some of the problems for Hawley's account: for example, the kind of counterexamples raised by Bengson & Moffett 2011, which seem to show that an agent can know how to ϕ even without being able to ϕ in any nearby world.

⁴²The 'C' appearing on the right-hand side of the definition, it is supposed to be a placeholder for some more detailed specification of relevant circumstances. For example, these circumstances might be the ones that are suitably close to the circumstances of c .

knew how to hit a home run in baseball and that Lionel Messi knows how to score a penalty in soccer. But there is a major difference in their success rate: Babe Ruth hit a home run in much less than 50% of the circumstances in which he was batting, while Messi's penalty scoring record is much higher. The fact that we still count Babe Ruth as possessing know-how is due to the fact that hitting a home run in baseball is much harder than scoring a penalty in soccer.

In summary, we can count a subject as having know-how in circumstances that vary widely along a number of dimensions. Here I don't need to settle what fixes the values of these parameters—whether features of the context, of the subject of know-how, or other elements. Similarly, I avoid commitments about whether and how know-how ascriptions end up being context-dependent in the relevant ways. These are important questions, but I don't have the space to address them here, and I must leave them to future work on the topic.⁴³

4.5 Factualism regained?

This concludes my proposed nonfactualist account of know-how. Before moving on, I want to consider a general line of criticism to the project. Quite simply, the objection is that the account is just a kind of factualism in disguise. At the very least, the account is much closer to factualism than Ryle's original account. I am claiming that states of know-how have content and that this content can be modeled by tools that are adapted from standard possible worlds semantics. I have gone far from Ryle, who identified know-how with an ability (and hence, given his view of abilities, with a set of dispositions).

I agree that I have come a long way from Ryle. But I have not gone all the way to the factualist side: I still retain a crucial point of disagreement with factualists. Recall Stanley's (2011) statement of factualism:

[K]nowing how to do something is the same as knowing a fact. It follows that learning how to do something is learning a fact.

This is exactly the main claim I deny. On the account I have developed, knowing how to do something involves entertaining a nonfactual kind of content. In turn, this difference is tied to a difference in functional roles between know-how and representational mental states.

Moreover, and more importantly, insisting on defending a purely Rylean position, at this stage of the debate, seems anachronistic. Using nonfactual contents seems the way forward if we aim to square nonfactualism with know-how ascriptions. Once more,

⁴³There are obvious analogies to the debate on contextualism about propositional knowledge here. For some relevant positions, see DeRose 1992, Lewis 1996, Hawthorne 2004, Stanley 2005; see also Rysiew 2011 for an overview.

the analogy with noncognitivism in metaethics is illuminating. Historical precursors of expressivism, for example the emotivism defended by Stevenson (1944), refrained from assigning anything like Gibbard-style content to normative claims. Rather, they took normative claims to be roughly on a par with exclamations such as ‘Boo!’ or ‘Hooray!’.⁴⁴ Views of this sort ran into the usual Frege-Geach worries. First, they were unable to account for the syntactic complexity of normative discourse. Second, they were unable to capture adequately the way in which statements involving both normative and descriptive discourse can stand in logical relations. Modern-day expressivism was developed, in part at least, as a response to these objections. The move that allowed the expressivist to avoid the linguistic and logical worries raised by emotivism was just allowing that normative claims express a specific, nonrepresentational kind of content. In essence, the move that I’m advocating mirrors the shift from archaic to contemporary forms of noncognitivism. The compositional worries raised by Stanley and Williamson can be met by granting that know-how has a kind of content. The nonrepresentational element in the account is retained via the claim that this content is different from the content of standard representational attitudes like beliefs.

Thus I agree that the account I have proposed is in several ways non-Rylean. But what I’m interested in is not vindicating Ryle, but rather developing a form of non-factualism that overcomes the objections raised against Ryle, while still retaining the main spirit of the position. If I’m right, the history of the debate about noncognitivism teaches us that this is the way to go.

5 Know-how reports: outline of a semantics

Stating a full semantics for know-how reports requires touching on technical issues. Hence the task is best left to an appendix. But here I can flag, in outline, what needs to change in standard accounts and how MPPs can be implemented into the semantics.

My starting point is Rajesh Bhatt’s semantics for infinitival questions (1999, 2006), which is the most developed existing work on the topic. On Bhatt’s account, all questions in the infinitival mood (such as *whom to invite to the party*, *what to cook for dinner*, and *how to cook risotto*) involve a special covert modal, which Bhatt represents as ‘ $\diamond_{D,\rightarrow}$ ’. This modal is a kind of bouletic modal: roughly, it concerns what a subject should do, in order to achieve certain goals. At the same time, differently from natural language *should*, ‘ $\diamond_{D,\rightarrow}$ ’ has a conjunctive meaning, involving both existential and universal quantification: ‘ $\diamond_{D,\rightarrow}(p)$ ’ says that *there are* some *p*-worlds where the relevant goals are satisfied, and that *all* the *p*-worlds are worlds where the goals are satisfied.

⁴⁴This quick summary obviously can’t do justice to the complexities of Stevenson’s argument. But my main focus is elsewhere and it will do for current purposes.

Why this double mechanism of quantification? Notice that knowledge reports involving infinitival questions can be paraphrased in different ways in different contexts. Consider (17):

(17) Sam knows what to tell his friends to make them happy.

Depending on context, the most appropriate paraphrase for (17) will be (18) or (19):

(18) Sam knows what he can tell his friends to make them happy.

(19) Sam knows what he should tell his friends to make them happy.

(18), the so-called ‘mention-some’ reading, involves an existential modal with circumstantial flavor (capturing what can happen, in view of the facts). (19), the so-called ‘mention-all’ reading, involves a universal modal with bouletic flavor (capturing what one should do, in view of certain goals one wants to achieve). Bhatt’s achievement consists in deriving both readings from only one meaning for the modal (plus some assumptions which will vary with context). To do this, he builds both quantifiers in the meaning of ‘ $\diamond_{D,\rightarrow}$ ’. In a slogan, ‘ $\diamond_{D,\rightarrow}$ ’ can be paraphrased, depending on the circumstances, as a *can* or as a *should* because its meaning just is the conjunction of a *can* and a *should*. (For the exact way in which the modal gives rise to the two readings, the reader is referred to Bhatt’s original discussion.)

Just the meaning of ‘ $\diamond_{D,\rightarrow}$ ’ is the natural point of intervention to operate the switch to a nonfactualist semantics. In particular, I modify the universal quantification element. Rather than universally quantifying over worlds, the modal ‘ $\diamond_{D,\rightarrow}$ ’ now quantifies universally over the courses of action that are ‘listed’ within MPPs.

Let me now introduce some basic assumptions about embedded questions. On standard semantics for knowledge reports involving questions, these reports state that the subject has knowledge of (some or all of) the true answers to the question. Similarly, on the view I suggest, a report of the form *S knows how to ϕ* states that *S* has knowledge of some of the answers that lead to successful completion of the relevant tasks, as spelled out in section 4.

Now, take my running example (1). Suppose, for simplicity, that there are three relevant ways to make risotto, and only one (label it ‘Way 1’) is conducive to Sam’s goal of making good quality risotto. The truth conditions we get for (1) in combination with Bhatt’s original semantics for the modal ‘ $\diamond_{D,\rightarrow}$ ’ are (roughly):

(1) is true iff, for all worlds w' compatible with Sam’s knowledge: (a) there is some world relevantly similar to w' where Sam makes good risotto by cooking risotto in Way 1 and (b) in all worlds relevantly similar to w' where Sam cooks risotto in Way 1, he cooks good risotto.⁴⁵

⁴⁵The notion of a ‘relevantly similar’ world is supposed to capture the flavor of so-called circumstantial

By contrast, the truth conditions we get on the nonfactualist semantics are (again, with some approximation):

(1) is true iff, for all world-MPP pairs $\langle w', m' \rangle$ compatible with Sam's knowledge: (a) there is some world relevantly similar to w' where Sam makes good risotto by cooking risotto in Way 1 and (b) *all courses of action in m' that involve cooking risotto in Way 1 are conducive to making good risotto, according to m' .*

The second clause, obviously, is where the difference lies. The reader is invited to consult the appendix for a more precise statement of the truth conditions, as well as for an explanation of how these truth conditions are derived compositionally.

6 Conclusion

The line defended by Stanley and Williamson has proved very influential in the literature on know-how, and rightly so. Stanley and Williamson point out an important flaw in classical versions of nonfactualism and they set a new standard that a theory of know-how should meet. But they are too quick in inferring, from this, the truth of factualism. We have the semantic resources to give a solid account of know-how that is able to reconcile standard semantic views with the claim that knowing how to do something doesn't consist in knowing a fact. Moreover, these semantic resources dovetail with a plausible view of the content of states of know-how in philosophy of mind. Know-how is a forward looking state, whose central functional role is guiding behavior and performance of tasks.

modality. This is the modality in play in statements like:

(20) Hydrangeas can grow around here.

Once more, the reader is referred to the appendix, as well as to Bhatt's discussion, for more details.

Appendix: nonfactualist semantics for know-how reports

To start with, a word of warning. The semantics of know-how reports lies at the intersection of a number of complex issues: the semantics of questions, the semantics of infinitival clauses, and the semantics of covert modality. While some important ground on this topic has been covered, much remains to be done. Hence it's unlikely that the account I'm about to suggest will be definitive. The following has mainly the purpose of showing how we can move to a nonfactualist framework starting from a fully standard factualist semantics.

1. Bhatt semantics for infinitival questions

It is broadly acknowledged that all embedded questions in the infinitival mood involve a kind of covert modality.⁴⁶ The easiest way to see this is to consider intuitive paraphrases of attitude ascriptions involving infinitival questions:

- (21) Sam knows whom to talk to about cooking.
- (22) Sam knows whom he/one should talk to about cooking.
- (23) Sam knows what to put in vegetable stock.
- (24) Sam knows what he/one should put in vegetable stock.
- (25) Sam knows where to find high-quality arborio rice.
- (26) Sam knows where he/one can find high-quality arborio rice.

Notice that different embedded questions are best paraphrased with different kinds of overt modals. In some cases (for example (21) and (23)), a necessity modal is most appropriate. This modal has a *bouletic* flavor: it concerns what a subject is required to do, given certain goals. In other cases (for example, (25)), a possibility modal is most appropriate. This modal has a *circumstantial* flavor, indicating what can be the case, given certain facts about the world.

Know-how reports follow this pattern. Most of them are most naturally paraphrased with *can*:

- (27) Sam knows how to ride a bike.
- (28) Sam knows how he/one can ride a bike.

⁴⁶In line with current syntactic literature (and in particular with the analysis in Bhatt 2006), I assume that the modal is realized as a covert element present in the syntactic structure of the sentence (in particular, as the complementizer C_o):

- (i) Sam_i knows how_i [C_o PRO₁ to cook risotto t_i]

Some other ones call for a *should*-paraphrase, especially in certain contexts. Suppose that we're discussing methods for cooking risotto and that I'm pointing out that performing certain operations is crucial for a good outcome. In this context, it seems appropriate to paraphrase (1) with (29):

- (1) Sam knows how to cook risotto.
 (29) Sam knows how he/one should cook risotto.

It's desirable to have an analysis of know-how reports that accommodates both the necessity and the possibility readings. My strategy will be this: I will use as my benchmark theory the account of infinitival modals developed by Rajesh Bhatt (1999 and 2006), and show how this account can be turned, via minimal tweaks, into a nonfactualist account. Bhatt's account succeeds in predicting the full range of data, while managing to assume a unitary meaning for the modal present in infinitival questions. The nonfactualist proposal I will develop will inherit this feature. Bhatt's semantics is a natural choice also for more general reasons. At the current state of play, it seems the best and most developed account of the semantics of infinitival question.⁴⁷

Bhatt's central idea is that all infinitival questions involve a special covert modal, which he represents as ' $\diamond_{D,\rightarrow}$ '. The meaning of this modal, on a par with standard bouletic modals, exploits reference to two sets of worlds. Both of them are determined as a function of an input world w . The first is a set of worlds that are relevantly similar to w . Following Bhatt, I refer to this set via the shorthand ' $Rel(w)$ '. The second is a set of worlds where certain relevant goals that the subject has in w are realized; following Bhatt again, I will refer to it via ' $Goal(w)$ '. This set will always be a subset of $Rel(w)$. $Rel(w)$ and $Goal(w)$ roughly coincide respectively with what are often called the *modal base* and the *ordering source*, in the terminology that is standard from the work of Kratzer (1981, 1991).

Now that we are equipped with this terminology, I can state the meaning of an arbitrary clause of the form ' $\diamond_{D,\rightarrow}(p)$ ' (where ' p ' stands for an infinitival clause that picks out a proposition). This meaning is conjunctive and it involves *both* existential and universal quantification over worlds:

$\llbracket \diamond_{D,\rightarrow}(p) \rrbracket^w$ is true iff both of the following conditions obtain:

- there is a world w' such that $w' \in p$ and $w' \in Goal(w)$;
- for all worlds w' in $Rel(w)$ where p is true (i.e. all worlds w' such that $w' \in p$), p brings it about that the relevant goals are satisfied in w' (i.e. that $w' \in Goal(w)$).

⁴⁷Let me emphasize that nothing in my general account relies on exploiting Bhatt's semantics. The nonfactualist strategy may be pursued in a number of different ways. Indeed, if my general line of argument in this paper is correct, virtually any factualist account may be turned into a nonfactualist one.

It's useful to give an intuitive gloss:

$\ulcorner \diamond_{D,\rightarrow} (p) \urcorner$ means: (a) that p is possible; and (b) that p invariably brings it about that a set of relevant goals are satisfied.

Finally, just for reference, here is a statement of the lexical entry for ' $\diamond_{D,\rightarrow}$ ':

$$\llbracket \diamond_{D,\rightarrow} \rrbracket^w = \lambda p. \exists w' [w' \in \text{Goal}(w) \text{ and } p(w')] \\ \text{and } \forall w' [\llbracket w' \in \text{Rel}(w) \text{ and } p(w') \rrbracket \rightarrow w' \in \text{Goal}(w)]$$

The double quantification in the meaning of ' $\diamond_{D,\rightarrow} (p)$ ' is exactly what makes the case that infinitival questions are sometimes heard as involving a possibility modal, and sometimes heard as involving a necessity modal. I skirt over the exact mechanics that determine one of the two readings; the reader is referred to Bhatt (2006, pages 129–142) for a full account.

It's also useful to go through a couple of examples. Consider first:

(21) Sam knows whom to talk to about cooking.

In order to interpret (21), we need to fix a set of goals that Sam has and that are relevant in the context. Assume that the only relevant goal for Sam is improving his cooking skills. Also, suppose that there are three relevant people in the domain of discourse: Pablo, Jane, and Max. Talking to Pablo and Jane would be conducive to an improvement in Sam's cooking skills, but not talking to Max.

As a next step, we need a semantics for questions: I choose Karttunen semantics (1977), on which an embedded question denotes a set of true propositions—intuitively, the set of its true answers. The denotation of the clause *whom to talk to about cooking* then is (somewhat simplifying) a set of two propositions:⁴⁸

$$\llbracket \text{whom}_1 \llbracket \diamond_{D,\rightarrow} [\text{PRO to talk to } t_1 \text{ about cooking}] \rrbracket \rrbracket^w = \\ \{p: p(w) = 1 \text{ and} \\ [p = \lambda w'. \diamond_{D,\rightarrow} (x \text{ talks to Pablo about cooking}) \text{ in } w', \text{ or} \\ p = \lambda w'. \diamond_{D,\rightarrow} (x \text{ talks to Jane about cooking}) \text{ in } w' \text{ or} \\ p = \lambda w'. \diamond_{D,\rightarrow} (x \text{ talks to Max about cooking}) \text{ in } w']\} = \\ \{\lambda w'. \diamond_{D,\rightarrow} (x \text{ talks to Pablo about cooking}) \text{ in } w', \\ \lambda w'. \diamond_{D,\rightarrow} (x \text{ talks to Jane about cooking}) \text{ in } w'\}$$

⁴⁸For simplicity, I'm treating the pronoun PRO as an ordinary variable and I'm not representing explicitly assignments in the semantics. For a proper semantics for PRO, see Chierchia 1989. Also, I'm just assuming that embedded wh-questions like *who to talk to about cooking* will make available a set of propositions at the right point in the tree. Since I'm using an intensional system, presumably this will require some rule in the style of the Intensional Functional Application rule stated by Heim & Kratzer 1998. Here I'm not worrying about working out the details of the semantics in a precise way. What I want to do is just to illustrate the main changes involved in going nonfactualist.

On Karttunen semantics for questions, a knowledge report involving an embedded infinitival question just means that the subject knows at least one of the propositions in the set.⁴⁹ The corresponding lexical entry for *know* is:

$$\llbracket \text{know} \rrbracket^w = \lambda x. \lambda S. \exists p \in S \text{ such that, for all } w' \text{ compatible with } x\text{'s knowledge in } w, w' \in p$$

In our case, (21) is predicted to mean that the subject knows one of the propositions in:

$$\{ \lambda w'. \diamond_{D,\rightarrow} (x \text{ talks to Pablo about cooking}) \text{ in } w', \\ \lambda w'. \diamond_{D,\rightarrow} (x \text{ talks to Jane about cooking}) \text{ in } w' \}$$

It's instructive to go through a simplified version of the compositional derivation. For simplicity, here I use the modal ' $\diamond_{D,\rightarrow}$ ' also in the metalanguage. I unpack the meaning of the modal (thus giving the full-blown truth conditions) below.

$$\begin{aligned} \llbracket \text{Sam knows who to talk to about cooking} \rrbracket^w = \\ & [\lambda S. \exists p \in S \text{ such that, for all worlds } w' \text{ compatible with Sam's knowledge in } w, p(w') = 1] (\llbracket \text{who to talk to about cooking} \rrbracket^w) = \\ & [\lambda S. \exists p \in S \text{ such that, for all worlds } w' \text{ compatible with Sam's knowledge in } w, p(w') = 1] (\{ \lambda w'. \diamond_{D,\rightarrow} (x \text{ talks to Pablo about cooking}) \text{ in } w', \\ & \lambda w'. \diamond_{D,\rightarrow} (x \text{ talks to Jane about cooking}) \text{ in } w' \}) = \\ & \exists p \in \{ \lambda w'. \diamond_{D,\rightarrow} (\text{Sam talks to Pablo about cooking}) \text{ in } w', \lambda w'. \diamond_{D,\rightarrow} (\text{Sam talks to Jane about cooking}) \text{ in } w' \} \text{ such that, for all worlds } w' \text{ compatible with Sam's knowledge in } w, p(w') = 1 \end{aligned}$$

Given Bhatt semantics for $\diamond_{D,\rightarrow}$, (21) gets the following truth conditions:

$$\llbracket (21) \rrbracket^w \text{ is true iff, for all } w' \text{ compatible with what Sam knows in } w, \text{ either:}$$

- there is a world w'' such that Sam talks to Pablo about cooking in w'' and Sam's goals in w' are satisfied in w'' (i.e. Sam improves his cooking skills in w'');
- for all worlds w'' that are relevantly similar to w' and in which Sam talks to Pablo about cooking, Sam's talking to Pablo about cooking brings it about that Sam's goals in w' are satisfied in w'' (i.e. that Sam improves his cooking skills in w'').

⁴⁹This is a simplification. In general, Karttunen's semantics requires that the subject knows *all* the true answers in the set (this is the so-called 'weak exhaustivity' requirement); but just this requirement is absent for the case of infinitival questions. To my knowledge, we don't have an account yet of what produces the failure of weak exhaustivity in these cases. Here I limit myself to assuming that this failure happens and that this is part of the meaning of the verb.

or:

- there is a world w'' such that Sam talks to Jane about cooking in w'' and Sam's goals in w' are satisfied in w'' (i.e. Sam improves his cooking skills in w'');
- for all worlds w'' that are relevantly similar to w' and in which Sam talks to Jane about cooking, Sam's talking to Jane about cooking brings it about that Sam's goals in w' are satisfied in w'' (i.e. that Sam improves his cooking skills in w'').

More simply, following the intuitive gloss: (21) is true iff John knows one of the following: (a) it's possible that he talks to Pablo about cooking and that talking to Pablo about cooking invariably brings it about that his cooking skills are improved; (b) it's possible that he talks to Jane about cooking and that talking to Jane about cooking invariably brings it about that his cooking skills are improved.

Since the truth conditions we get are pretty complex, from now on I'll just consider examples where the relevant wh-complement denotes a singleton—i.e. examples where the relevant question has only one true answer. This will simplify things and allow us to focus on the key innovation of Bhatt semantics, i.e. the semantics of the modal $\diamond_{D,\rightarrow}$.

Take then our running example of a know-how report:

- (1) Sam knows how to cook risotto.

To fix a context, suppose that Sam has one relevant goal: making good risotto. To simplify, suppose also that there are three relevant ways, Way 1, Way 2, and Way 3, and that Way 1 is the only one conducive to making good risotto. The denotation of the complement clause in (1) is then:

$$\llbracket \text{how to cook risotto} \rrbracket^w = \{ \lambda w'. \diamond_{D,\rightarrow} (x \text{ cooks risotto in Way 1}) \text{ in } w' \}$$

The truth conditions assigned to (1) by Karttunen semantics for questions combined with the Bhatt machinery are the following:

$\llbracket (1) \rrbracket^w$ is true iff, for all w' compatible with what Sam knows in w ,

- there is a world w'' such that Sam cooks risotto in Way 1 in w'' and Sam's goals in w' are satisfied in w'' (i.e. Sam makes good risotto in w'');
- for all worlds w'' that are relevantly similar to w' and in which Sam cooks risotto in Way 1, Sam's cooking risotto in Way 1 brings it about that Sam's goals in w' are satisfied in w'' (i.e. that Sam makes good risotto in w'').

More simply, (1) is true just in case Sam knows that cooking risotto in Way 1 is possible, and that cooking risotto in Way 1 invariably brings it about that he makes good risotto.

II. Going nonfactualist

Bhatt’s account has factualist assumptions in the background. In line with standard accounts of embedded questions, the knowledge ascribed by a know-how report is assumed to be knowledge of a proposition. But this assumption can be replaced. We can swap it for a nonfactualist assumption with only minimal adjustments. Below, I will demonstrate how to get a nonfactualist semantics for all knowledge reports involving infinitival embedded questions. I’m going to briefly discuss in the next section whether and to what extent this approach should be restricted to know-how reports.

As I highlighted in section 4, one first key nonfactualist maneuver is a switch in the basic atoms of possible worlds semantics. Bhatt simply assumes that *know* and other attitude verbs taking infinitival embedded questions as complements quantify over worlds. I take them rather to quantify over pairs of a world and an MPP. As a result, exactly as it happens in Gibbard’s semantics for normative language, a know-how report may ascribe knowledge of a plan rather than of a proposition.⁵⁰ Here is the new semantics for *know*:

$$\llbracket \text{know} \rrbracket^{w,m} = \lambda x. \lambda S. \exists p \in S \text{ such that, for all } \langle w', m' \rangle \text{ compatible with } x\text{'s} \\ \text{knowledge in } w, \langle w', m' \rangle \in p$$

Notice that the variable ‘*m*’ is not mentioned on the right-hand side. This shows that the denotation of *know* is not sensitive to the parameter tracking MPPs. This is as it should be: it is a factual matter what propositional knowledge or know-how a subject has. At the same time, *know* will shift the MPP parameter at which its complement gets evaluated. As we will see in a minute, the MPP parameter there is *not* going to be idle.

As I mentioned in section 4, the notion of an MPP being compatible with a subject’s knowledge will be explicated in terms of the notion of an MPP being compatible with a subject’s plan, which is the primitive of the theory. As for the notion of a world-MPP pair being compatible with a subject’s knowledge, I understand it, in the obvious way, as the world and the MPP being both compatible with what the subject knows.

The second major change concerns the covert modality investigated by Bhatt. On Bhatt semantics, the covert modal took as input knowledge-worlds—i.e., worlds compatible with the subject’s knowledge—and used them in various ways (for example, to determine the sets *Rel*(*w*) and *Goal*(*w*)). Schematically:

⁵⁰In particular, which state is ascribed is going to be determined by the kind of complement involved. As I’m going to point out below, I assume that the covert modal $\diamond_{D,\rightarrow}$ looks at the MPP argument rather than at the world one, and hence it forces ascription of a plan rather than of a proposition.

$\llbracket S \text{ knows } \dots \diamond_{D,\rightarrow} \dots \rrbracket = \text{true iff, for every knowledge world } w' \dots \text{ for all worlds } w'' \text{ relevantly similar to } w' \text{ and for all worlds } w'' \text{ where the subject's goals in } w' \text{ are satisfied, } \dots$

The result of this is that the question embedded under *know* end up saying something about the knowledge-worlds of the subject. It is this piece of the account that we want to change. So I substitute MPPs for worlds. The modal $\diamond_{D,\rightarrow}$ will now take as arguments MPPs and the embedded question will say something about the subject's plans. This move, combined with the switch in the denotation of *know*, is what generates the switch to nonfactualism.

This change calls for a further adjustment. On Bhatt semantics, each knowledge world is used to determine two sets of worlds $Rel(w)$ and $Goal(w)$ —a set of relevantly similar worlds and a set of goal-worlds. On my semantics, we still use knowledge-worlds to determine a set of contextually relevant goal-worlds $Goal(w)$. But, rather than the set $Rel(w)$, we have a different parameter: a set of sets of worlds, representing *a set of possible courses of action*. (Recall from section 4 that a course of action can be modeled just as a set of worlds.) This parameter is going to be determined as a function of the MPP parameter, rather than of the world parameter. I represent it as ' $Act(m)$ '.⁵¹

At this point, I am able to state the new semantics for the modal $\diamond_{D,\rightarrow}$. Informally, and with some approximation, what the modal does is the following: it looks inside an MPP m , and says that, according to m , certain courses of action bring about certain goals. Here is the new meaning:

$\llbracket \diamond_{D,\rightarrow} (p) \rrbracket^{w,m}$ is true iff both of the following conditions obtain:

- there is a world w' such that $w' \in p$ and $w' \in Goal(w)$;
- according to m , for all courses of actions A in $Act(m)$ such that $A \subseteq p$ (i.e. for all courses of action that include action p), p brings it about that the relevant goals are satisfied throughout A (i.e. that $A \subseteq Goal(w)$)

Here is the intuitive gloss on the new semantics: ' $\diamond_{D,\rightarrow} (p)$ ' is true, relative to a world w and an MPP m , just in case: (a) p is contemplated by some successful courses of action; and (b) according to m , courses of action involving p invariably bring it about that the contextually relevant goals are satisfied. The new lexical entry is the following:

$$\llbracket \diamond_{D,\rightarrow} \rrbracket^{w,m} = \lambda p. \exists w' [w' \in Goal(w) \text{ and } p(w')] \\ \text{and } \forall A [[A \in Act(m) \text{ and } A \subseteq p(w')] \rightarrow A \subseteq Goal(w) \text{ in } m]$$

⁵¹If we take all MPPs to map all possible courses of actions into goals, then $Act(m)$ will be the same for all MPPs and will just be the set of all possible courses of action. But I think it's useful to give a statement of the semantics that is independent of this assumption.

Notice that, on the new semantics, the modal doesn't quantify over MPPs any more, but rather it quantifies (universally) over the courses of actions listed by an MPP. This is the major shift in the whole compositional machinery. So far as I can see, it doesn't affect any of the compositional properties of ' $\diamond_{D,\rightarrow}$ '.⁵² Notice that, in contexts where the modal ' $\diamond_{D,\rightarrow}$ ' does appear embedded under an attitude verb like *know*, the truth conditions of the whole sentence *will* still involve universal quantification over MPPs, thanks to the quantificational force of the attitude verb. What we don't have is quantification over MPPs twice over, as it happened in Bhatt's original entry for the case of worlds.⁵³

This new meaning of the modal produces a new meaning for infinitival questions. It's useful to give an example. Suppose, similarly to what we did above, that there is only one set of instructions that reliably guides a subject to cooking risotto successfully, i.e. cooking risotto in Way 1. We get that the denotation of *how to cook risotto* is:

$$\begin{aligned} \llbracket \text{how to cook risotto} \rrbracket^{w,m} = \\ \{ \lambda \langle w', m' \rangle. \diamond_{D,\rightarrow} (x \text{ cooks risotto in Way 1}) \text{ at } \langle w', m' \rangle \} = \\ \{ \lambda \langle w', m' \rangle. \text{There is a world } w'' \text{ such that } x \text{ cooks risotto in Way 1 in } w'' \\ \text{and the goals in } \textit{Goal}(w') \text{ are satisfied in } w''; \text{ and, for all courses of actions} \\ A \text{ in } \textit{Act}(m') \text{ such that } A \text{ includes } x \text{ cooking risotto in Way 1, } x \text{ cooking} \\ \text{risotto in Way 1 brings it about that the goals in } \textit{Goal}(w') \text{ are satisfied} \} \end{aligned}$$

The compositional derivation is essentially analogous to the one given for the factulist case, so I skip it. Let me just state the truth conditions of my running example (1). Again, I assume that there are only three ways to make risotto, and that the only goal is making good risotto, and that the only set of instructions conducive to success is the one that implements Way 1. We get:

⁵²Let me also notice that the idea of quantifying over courses of actions is echoed by some recent and less recent work on the semantics of deontic modals (see e.g. Jackson & Pargetter 1986, Cariani 2013).

⁵³The asymmetry can be eliminated if we move away from the standard way of construing the interaction between attitude verbs and modals and adopt a new picture inspired by recent work in semantics, for example by Yalcin 2007 and Gillies 2010. The basic idea is that when we have a configuration of the form

VERB . . . [. . . MODAL . . .]

the attitude verb shifts directly the domain of quantification of the embedded modal via a domain parameter. For example, the semantic effects of *know* can be represented schematically as follows:

$$\llbracket \text{S knows } [\phi] \rrbracket^{w,S} = \text{For all } w' \text{ compatible with } K_S, \llbracket \phi \rrbracket^{w',K_S}$$

where ' K_S ' represents S 's knowledge state. Correspondingly, modals quantify over the domain specified by the new parameter; for example, schematically again:

$$\llbracket \text{must } \phi \rrbracket^{w,S} = \forall w' \in S, \llbracket \phi \rrbracket^{w',S} = 1.$$

If we construe the interaction between attitude verbs and the special modal $\diamond_{D,\rightarrow}$ on these lines, we *can* build quantification over MPPs directly into the meaning of the modal. As I said in the main text, it's not clear to me that this makes any real difference to the computation of truth conditions, so I don't develop this line of thought explicitly here.

$\llbracket(1)\rrbracket^{w,m}$ is true iff, for all world-MPP pairs $\langle w', m' \rangle$ compatible with what Sam knows in w ,

- there is a world w'' such that Sam cooks risotto in Way 1 in w'' and Sam's goals in w' are satisfied in w'' (i.e. Sam makes good risotto in w'');
- according to m' , for all courses of actions A in $Act(m')$ such that Sam cooks risotto in Way 1 in A , Sam's cooking risotto in Way 1 brings it about that Sam's goals in w' are satisfied (i.e. that Sam makes good risotto) in A .

The intuitive gloss is now this: (1) is true just in case Sam knows that cooking risotto in Way 1 is possible; and, according to all MPPs compatible with Sam's knowledge, all courses of action that involve cooking risotto in Way 1 bring it about that Sam makes good risotto.

Let me close by summarizing the main moves of this section. the switch to a nonfactualist semantics relies on two main alterations of Bhatt semantics. First, *know* quantifies over pairs of a world and an MPP, rather than just a world. Second, the modal $\diamond_{D,\rightarrow}$ looks at all the subject's MPPs and checks that certain connections hold within them. In particular, it checks that p -courses of action are courses of action conducive to the achievement of goals. This mirrors very closely what happened in the original, factualist version. But it avoids the commitment to factualism.

III. Final comment: a generality problem?

Before closing, let me briefly consider an objection. I have specified a semantics that can assign nonfactualist truth conditions to all attitude reports involving infinitival questions. One might introduce extra stipulations: for example, one may claim that ' $\diamond_{D,\rightarrow}$ ' is ambiguous and that the nonfactualist version only appears in know-how reports. But unless one does something of this sort, We will assign truth conditions involving reference to nonfactualist mental states also to other knowledge-wh reports with infinitival complements, such as *John knows whom to invite for dinner*. Hence the scope of application of my proposed semantics is much wider than know-how reports. Is this a problem and, if so, does it disqualify my account?

One quick answer is that the present account is no worse off than any other existing account on the market. Any account of the semantics of know-how reports that takes compositionality seriously will run into the problem of setting know-how reports aside from reports involving other infinitival questions. On the upside, solutions to the problem can be found rather easily; on the downside, they invariably have a stipulative flavor. For example, Stanley and Williamson must stipulate that the practical modes

of presentation that, according to them, figure in know-how reports may not figure in other kinds of embedded questions. This assumption comes with no independent justification and seems to impose an arbitrary constraint on the distribution of a linguistic element. Hence it seems a rather *ad hoc* stipulation. A stipulation with the same effects and similar theoretical costs, *mutatis mutandis*, could be imposed on my account.

One other possibility, of course, is that we decide that we should have a nonfactualist account for all infinitival embedded questions. This seems an appealing route to go for several reasons: aside from theoretical unity, it would allow us to single out one kind of linguistic construction (and in particular, one kind of modality) that has the job of describing forward-looking mental states. At the moment this is just a suggestive option, but it clearly deserves serious investigation.

References

- Adams, Marcus (2009). "Empirical Evidence and the Knowledge-That/Knowledge-How Distinction." *Synthese*, 170(1): pp. 97–114.
- Bengson, John, and Marc A. Moffett (2011). "Non-Propositional Intellectualism." In J. Bengson, and M. A. Moffett (eds.) *Knowing How: Essays on Knowledge, Mind, and Action*, Oxford University Press, pp. 161–195.
- Bengson, John, Marc A. Moffett, and Jennifer C. Wright (2009). "The Folk on Knowing How." *Philosophical Studies*, 142(3): pp. 387–401.
- Bhatt, Rajesh (1999). *Covert Modality in Non-finite Contexts*. Ph.D. thesis, University of Pennsylvania.
- Bhatt, Rajesh (2006). *Covert Modality in Non-finite Contexts*. Mouton de Gruyter.
- Cariani, Fabrizio (2013). "Ought and Resolution Semantics." *Noûs*, 47(3): pp. 534–558.
- Carr, David (1981). "Knowledge in Practice." *American Philosophical Quarterly*, 18: pp. 53–61.
- Cath, Yuri (2012). "Knowing How without Knowing That." In J. Bengson, and M. Moffett (eds.) *Knowing How: Essays on Knowledge, Mind, and Action*, Oxford University Press.
- Charlow, Nate (2013). "Logic and Semantics for Imperatives." *Journal of Philosophical Logic*, pp. 1–48.
- Charlow, Nate (2014). "The Meaning of Imperatives." *Philosophy Compass*, 9(8): pp. 540–555.
- Chierchia, Gennaro (1989). "Anaphora and Attitudes *De Se*." In R. Bartsch, J. van Benthem, and P. van Emde Boas (eds.) *Semantics and Contextual Expression*, Dordrecht: Foris.
- Cohen, Neal J., and Larry R. Squire (1980). "Preserved Learning and Retention of Pattern-Analyzing Skill in Amnesia: Dissociation of Knowing How and Knowing That." *Science*, 210: pp. 10.
- DeRose, Keith (1992). "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research*, 52(4): pp. 913–929.
- Dretske, Fred I. (1981). *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.

- Dummett, Michael (1981). *Frege: Philosophy of Language (second edition)*. Cambridge, MA: Harvard University Press.
- Gibbard, Allan (1990). *Wise choices, apt feelings: A theory of normative judgment*. Cambridge MA: Harvard University Press.
- Gibbard, Allan (2003). *Thinking How to Live*. Cambridge MA: Harvard University Press.
- Gillies, Anthony S (2010). "Iffiness." *Semantics and Pragmatics*, 3.
- Ginet, Carl (1975). *Knowledge, Perception and Memory*, vol. 5 of *Philosophical Studies Series in Philosophy*. Dordrecht, Holland: D. Reidel.
- Glick, Ephraim (2011). "Two Methodologies for Evaluating Intellectualism." *Philosophy and Phenomenological Research*, 82(1): pp. 1–37.
- Hawley, Katherine (2003). "Success and Knowledge-How." *American Philosophical Quarterly*, 40(1): pp. 19–31.
- Hawthorne, John (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Heim, Irene, and Angelika Kratzer (1998). *Semantics in Generative Grammar*. Oxford and Malden, MA: Blackwell Publishers.
- Jackson, Frank, and Robert Pargetter (1986). "Oughts, Options, and Actualism." *The Philosophical Review*, 95(2): pp. 233–255.
- Kaplan, David (1989a). "Demonstratives." In J. Almog, J. Perry, and H. Wettstein (eds.) *Themes from Kaplan*, Oxford: OUP.
- Kaplan, David (1989b). "Demonstratives: Afterthoughts." In J. Almog, J. Perry, and H. Wettstein (eds.) *Themes from Kaplan*, Oxford: OUP.
- Karttunen, L. (1977). "Syntax and semantics of questions." *Linguistics and philosophy*, 1(1): pp. 3–44.
- King, Jeffrey (2003). "Tense, Modality, and Semantic Values." *Philosophical Perspectives*, 17(1): pp. 195–246.
- Kratzer, Angelika (1981). "The Notional Category of Modality." In H. J. Eikmeyer, and H. Rieser (eds.) *Words, Worlds, and Contexts: New Approaches to Word Semantics*, Berlin: de Gruyter.

- Kratzer, Angelika (1991). "Conditionals." In A. V. Stechow, and D. Wunderlich (eds.) *Semantik: ein internationales Handbuch der zeitgenössischen Forschung*, Amsterdam: de Gruyter.
- Kratzer, Angelika (2012). *Modals and conditionals: New and revised perspectives*, vol. 36. Oxford University Press.
- Kusumoto, Kiyomi (2005). "On the Quantification over Time in Natural Language." *Natural Language Semantics*, 13(4): pp. 317–357.
- Lewis, David K. (1974). "Radical Interpretation." *Synthese*, 27(3): pp. 331–344.
- Lewis, David K. (1979). "Attitudes *De Dicto* and *De Se*." *Philosophical Review*, 88(4): pp. 513–43.
- Lewis, David K. (1980). "Index, Context and Content." In *Papers in Philosophical Logic*, Oxford: Oxford University Press.
- Lewis, David K. (1996). "Elusive Knowledge." *Australasian Journal of Philosophy*, 74(4): pp. 549–567.
- MacFarlane, John (2005). "Making Sense of Relative Truth." *Proceedings of the Aristotelian Society*, 105: pp. 321–39.
- MacFarlane, John (2014). *Assessment sensitivity: Relative truth and its applications*. Oxford University Press.
- Moss, Sarah (2013). "Epistemology Formalized." *Philosophical Review*, 122(1): pp. 1–43.
- Moss, Sarah (2014). "On the Semantics and Pragmatics of Epistemic Vocabulary." Forthcoming in *Semantics and Pragmatics*.
- Ninan, Dilip (2010). "Semantics and the Objects of Assertion." *Linguistics and Philosophy*, 33(5): pp. 355–380.
- Noë, Alva (2005). "Against Intellectualism." *Analysis*, 65(288): pp. 278–290.
- Percus, Orin (2000). "Constraints on some other variables in syntax." *Natural Language Semantics*, 8(3): pp. 173–229.
- Pérez Carballo, Alejandro (2012). "Negation, Expressivism, and Intentionality." Manuscript, University of Southern California.
- Perry, John (1979). "The Problem of the Essential Indexical." *Noûs*, 13: pp. 3–21.

- Pettit, Dean (2002). "Why knowledge is unnecessary for understanding language." *Mind*, 111(443): pp. 519–550.
- Rabern, Brian (2012). "Against the Identification of Assertoric Content with Compositional Value." Forthcoming in *Synthese*.
- Rothschild, Daniel (2012). "Expressing Credences." In *Proceedings of the Aristotelian Society*, vol. 112, pp. 99–114. Wiley Online Library.
- Ryle, Gilbert (1949). *The Concept of Mind*. Chicago IL: The University of Chicago Press.
- Rysiew, Patrick (2011). "Epistemic Contextualism." In E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy (Winter 2011 Edition)*, URL = <<http://plato.stanford.edu/archives/win2011/entries/contextualism-epistemology/>>.
- Schroeder, Mark (2008a). *Being For: Evaluating the Semantic Program of Expressivism: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford University Press.
- Schroeder, Mark (2008b). "How Expressivists Can and Should Solve Their Problem with Negation." *Noûs*, 42(4): pp. 573–599.
- Schroeder, Mark (2009). *Noncognitivism in Ethics*. Routledge.
- Stalnaker, Robert (1984). *Inquiry*. Cambridge, MA: MIT Press.
- Stanley, Jason (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Stanley, Jason (2011). *Know How*. Oxford: Oxford University Press.
- Stanley, Jason, and Timothy Williamson (2001). "Knowing how." *Journal of Philosophy*, 97: pp. 411–44.
- Stevenson, Charles (1944). *Ethics and Language*. New Haven: Yale University Press.
- Swanson, Eric (2006). *Interactions with Context*. Ph.D. thesis, MIT.
- Swanson, Eric (2011). "How not to Theorize about the Language of Subjective Uncertainty." In A. Egan, and B. Weatherson (eds.) *Epistemic Modality*, Oxford: Oxford University Press, pp. 249–269.
- von Fintel, Kai, and Irene Heim (2011). "Notes on Intensional Semantics." Unpublished lecture notes, MIT.

- Wallis, Charles (2008). "Consciousness, Context, and Know-how." *Synthese*, 160(1): pp. 123–153.
- Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Winograd, Terry (1975). "Frame Representations and the Declarative/Procedural Controversy." *Representation and Understanding: Studies in Cognitive Science*, pp. 185–210.
- Yalcin, Seth (2007). "Epistemic Modals." *Mind*, 116(464): pp. 983–1026.
- Yalcin, Seth (2011). "Nonfactualism about Epistemic Modality." In A. Egan, and B. Weatherson (eds.) *Epistemic Modality*, Oxford: Oxford University Press, pp. 295–332.
- Yalcin, Seth (2012a). "Bayesian Expressivism." In *Proceedings of the Aristotelian Society*, vol. 112, pp. 123–160. Wiley Online Library.
- Yalcin, Seth (2012b). "Context Probabilism." In *Proceedings of the 18th Amsterdam Colloquium*, pp. 12–21. Institute for Logic, Language, and Computation.