

Dynamic “Might” and Correct Belief*

Patrick Skeels

University of California, Davis

Synopsis: Veltman’s *test semantics* and developments thereof reject the canon about semantic contents and attitude ascriptions in favor of dynamic alternatives.¹ According to these theories the semantic content of a sentence is not a proposition, but a context change potential (CCP). Similarly, beliefs are not taken to be relations between agents and propositions, but agents and CCPs. These deviations from the canon come at the cost of an elegant explanation about the correctness of belief. Standardly, it is taken that the content of a belief is correct just in case the content of that belief is true. The proponent of the test semantics cannot appeal to this explanation since they hold that certain contents, namely epistemically modalized contents, do not express propositions, and are neither true nor false. Willer (2013) concerns how the test semantics can be marshalled to solve *inter alia* puzzles of modal disagreement. Crucial to Willer’s account is the proposal of a correctness condition which I call *evidential correctness*. According to evidential correctness, the content of a belief is correct just in case someone with the believing agent’s evidence would not be committed to factual error. The present paper argues that evidential correctness does not yield the proper correctness judgments for interlocutors in several common instances of modal disagreement. It further argues that evidential correctness concerns a different kind of correctness than the kind operative in modal disagreements. Since correctness judgments are what motivate the claim that modal disagreements are genuine disagreements, I take these objections to be significant. I subsequently consider two potential alternatives: *weak correctness* and *strong correctness* but conclude that each suffers from substantial problems. I then provide my own novel correctness conditions for belief contents within a dynamic framework. I argue that there are, in fact, two conditions for correctness of content that ought to be considered when operating within the test semantics. The first, *locative correctness*, applies only to contents which locate the believing agent in some subset of the space of possibility by entailing some contingent proposition. The second, *informational correctness* applies to contents that do not locate the believing agent. Such content includes, most notably, epistemically modalized content. After motivating this distinction, I

*Forthcoming in *Inquiry: an Interdisciplinary Journal of Philosophy*.

¹The test semantics is introduced in Veltman (1996). See Groenendijk and Stokhof (1991a), Groenendijk and Stokhof (1991b), and van der Does et al. (1997) among many, many others for further developments.

demonstrate how it yields the requisite correctness judgments on the part of disagreeing agents, and avoids the problems of the previously considered views.

Introduction

Canonical accounts of semantic content hold that the semantic value of a sentence is a proposition. Canonical accounts of belief hold that belief is a relation between an agent and a proposition. Thus, the canon often helps itself to an elegant explanation about the correctness of belief: a belief is correct if and only if the content of that belief is true, (Wedgewood, 2002, 267). Hence the slogan “belief aims at truth.”² A popular, dynamic, story about epistemic modals, often called *update semantics* or *test semantics*, rejects the canon about content.³ Introduced in Veltman (1996), and developed extensively thereafter, the test semantics has it that contents are not static propositions which are true or false. Instead, they are dynamic updates—functions from contexts to contexts—that provide instructions on how to update a body of information. Thus, meanings are not propositions but context change potentials. Propositions nonetheless play a fundamental role in theorizing and are used to model information states and truth. In addition, many updates can be said to express propositions, however, certain epistemically modalized updates do not (Veltman, 1996, 231). Instead, they perform the eponymous test on an information state. Let an information state be a set of possible worlds compatible with some body of information. Updating with “It might be the case that p ” tests the state for compatibility with p . If the state contains any worlds where p is true, then the initial state is returned (the test is passed). If there are no p worlds then \emptyset is the output (the test is failed), (Veltman, 1996, 228). Importantly, agents can have epistemically modalized beliefs, and if the contents of these beliefs are not propositions, then belief cannot be characterized as a relation between an agent and a proposition. Thus, proponents of the test semantics also reject the canon about belief. Instead, dynamic accounts often adopt an alternative defended in Heim (1992), according to which, belief is defined as the fixed point of an update. An agent believes some content just in case their information state remains unchanged when updated with that content. The relevant advantage of this view of belief is that the associated content need not be a proposition, and thus, agents can hold non-propositional beliefs.

So far, so good, but rejecting the canons comes at an indirect cost. Notably, the proponent of the test semantics cannot appeal to the elegant explanation about correct belief mentioned above. Since it is held that not all beliefs are

²See Williams (1973) for the earliest instance of this slogan.

³See Groenendijk and Stokhof (1991a) van der Does et al. (1997) Beaver (2001), Aloni (2001), Gillies (2004), von Stechow and Gillies (2007), Dever (2013), Yalcin (2012b), Yalcin (2015), Willer (2013), and Lennertz (2019) among others, for developments and applications of the test semantics. For influential dynamic accounts that do not directly concern epistemic modals, see Kamp (1981), Heim (1983), and Groenendijk and Stokhof (1991b). See Goldstein (2019) for an update semantics where not all instances of the epistemic modal operator perform a test.

true or false, their correctness cannot be explained exclusively in terms of truth and falsity.⁴ Instead, dynamic theorists appeal to correctness conditions that are related to truth only indirectly. Willer (2013) provides what is, to my knowledge, the only dedicated account of correctness of belief for the test semantics. Call this *evidential correctness*. According to evidential correctness, a belief is correct if and only if it would not commit someone with the believing agent’s evidence to factual error, (Willer, 2013, 64). Crucially, this proposal does not require that the belief itself be true or false.

Willer (2013) develops the test semantics in order to solve puzzles of *modal disagreement*. Modal disagreements paradigmatically involve an agent *a* who believes that some proposition *p* is not the case and some agent *b* who believes that *p* might be the case. It is claimed that *a* disagrees with *b*, and the fact that the test semantics treats these contents as inconsistent is touted as a motivation to adopt the view. The puzzle is often marshalled against the contextualist semantics for epistemic modals because contextualism allegedly fails to explain the disagreement. It is further argued that modal disagreements are genuine disagreements *because a* is disposed to judge the content of *b*’s belief as wrong, incorrect, or mistaken.⁵ This is supported by the observation that *a* can respond with something like, “No, you’re wrong,” or “No, that’s not right,” in response to *b*’s assertion.⁶ This is to say that *a* is disposed to judge *b*’s belief as incorrect, and that this judgement supports the claim that *a* disagrees with *b*.

The present paper argues that evidential correctness fails to yield the proper correctness judgments in a wide variety of cases. In short, if *a* does not share evidence with *b*, and *a* makes correctness judgments based on evidential correctness, then there will be cases where *a* will not conclude that the content of *b*’s belief is incorrect since *b*’s belief that it might be that *p* would not commit anyone with *b*’s evidence to factual error. *a*, nonetheless, disagrees with *b*. Unfortunately, many paradigmatic modal disagreement scenarios concern cases precisely like this, and evidential correctness fails to provide the motivating normative judgment on the part of *a*. This is a problem, since the puzzle of modal disagreement is, itself, motivated by the claim that *a* judges that the content of *b*’s belief is incorrect. Since these correctness judgments motivate the problem of modal disagreement, proponents of the test semantics must provide a tenable alternative that yields the appropriate judgment, i.e. *a* must consider the content of *b*’s belief to be incorrect in all cases of modal disagreement. Otherwise, the disagreement challenge loses its teeth. I further argue evidential correctness concerns the wrong kind of correctness required to motivate the challenge. I conclude that we are forced to explore alternatives.

I next consider two possible alternatives for correctness of belief: *weak correctness* and *strong correctness*. I conclude, however, that despite their advantages, both are implausible. I then propose a distinction between two kinds of

⁴This differentiates dynamic semantics from several other popular theories, including contextualism and relativism, since both theories still take meaning to be truth-conditional.

⁵Henceforth, I shall use normative terms like “wrong”, “incorrect”, etc. interchangeably.

⁶Appeal to correctness judgments that motivate disagreement can be found in MacFarlane (2011), Willer (2013), and Lennertz (2019).

correctness that govern the different types of content within the test semantics. Call these two kinds of correctness *locative correctness* and *informational correctness*. In brief, content that accurately individuates the way the world is from the way that it is not, is locatively correct. This is to say that such content locates the believer within a contingent set of worlds which contains the actual world. Similarly, content that misrepresents the way the world is is locatively incorrect. However, beliefs that do not locate the believer in one world or another, e.g. “might” beliefs, neither represent nor misrepresent the world. Thus, they are neither correct nor incorrect, locatively speaking. Such contents, I argue, are instead beholden to the norm of *informational correctness*. Some content is informationally correct just in case that content is consistent with the information that represents the facts. Thus, instead of judging the correctness of epistemically modalized contents with respect to their capacity to represent (which they cannot do), they are judged by their capacity to guide inquiry towards the truth (which they can do). The present paper develops these notions in a way that is consistent with the correctness judgments that motivate the test semantics. In addition, I briefly explain why this doxastic framework is not only plausible, but preferable, in that it distinguishes between the manner in which not believing truly fails to be correct, and the manner in which believing falsely fails to be correct. I then show that the framework makes plausible predictions with respect to how we ought to believe under uncertainty, and how we utilize doxastic norms to guide our doxastic actions.

§1 introduces the puzzle of modal disagreement and briefly rehearses the test semantics. §2 develops evidential correctness, weak correctness, and strong correctness. Problems are posed for each, and none are considered satisfactory. §3 considers some possible lines of response to the challenges posed in §2 but concludes that none are successful. §4 presents and explains locative and informational correctness then demonstrates how adopting both avoids the problems of weak, strong, and evidential correctness. §5 closes with some brief remarks.

1 Preliminaries

1.1 Modal Disagreement

Consider the following well-rehearsed scenario: Holmes believes that Moriarty is not the murderer ($\neg m$). Watson is undecided, and believes both that Moriarty might be the murderer (Might m) and that he might not be (Might $\neg m$).⁷ Numerous criticisms of the contextualist semantics for epistemic modals rest on the claim that in scenarios like this, Holmes *disagrees* with Watson.⁸ Call such cases *modal disagreements*.⁹ The idea is that disagreement manifests in

⁷Here, and henceforth, the modal auxiliary “might” is to be interpreted epistemically.

⁸See [Kratzer \(1977\)](#) and [Kratzer \(1981\)](#) for characterizations of contextualism.

⁹Several variations of this scenario exist. This version is structurally identical to the one from [Lennertz \(2019\)](#) which characterizes disagreement as a relation between two believing agents that holds in virtue of their belief states. Upon this construal, disagreement is treated as a state, rather than an action, and disagreements need not manifest directly in conversation.

cases where the disagreeing agents hold beliefs that are inconsistent with each other’s. According to contextualism, $\neg m$ and *Might* m are consistent, and the disagreement goes unexplained.¹⁰ Crucially, the pressure to consider such cases as genuine disagreements is often motivated by the observation that Holmes believes, and is licensed to assert, something like “No, you’re wrong,” directed at Watson’s belief (or assertion) that Moriarty might be the murderer.¹¹ Similar scenarios involving diachronic intrapersonal disagreement have also been considered, where an agent who at one point believed that p might be the case, and later learns that p is not the case, can remark “I used to believe that p might be the case, but I was wrong.”¹² In each case, some agent holds a belief, and in virtue of holding that belief, they make a normative judgment that some other belief (and the agent who holds it) is wrong. This judgment is treated as evidence that the first agent disagrees with the second. Contextualism’s alleged failure to properly predict these instances of disagreement is often used to motivate various non-canonical alternatives, including the test semantics.

1.2 Test Semantics

We now consider the test semantics in a simple and fairly general form. More sophisticated variants, most notably the semantics in [Willer \(2013\)](#), will be discussed in §3.¹³

Let \mathcal{L}_1 be generated by the grammar:

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid \text{Might } \varphi$$

Definition 1: Update System

$\langle \mathcal{L}, C, \cdot[\cdot] \rangle$ is an update system if and only if \mathcal{L} is a set of sentences, C is a set of information states, and $\cdot[\cdot]$ is a function which maps sentences of \mathcal{L} to operations on C .

Let W be the set of all functions $\mathcal{A} \mapsto \{0, 1\}$. Our update system will be $\langle \mathcal{L}_1, \mathcal{P}(W), \cdot[\cdot] \rangle$. If s is a state (a set of worlds in $\mathcal{P}(W)$) and φ is formula of \mathcal{L}_1 , then $s[\varphi]$ is the result of updating s with a formula φ . We now define the relevant notions of support and entailment, in addition to the semantic clauses for the connectives.

Definition 2: Support

$$s \models \varphi \Leftrightarrow s[\varphi] = s$$

For variants of the puzzle, see [Egan et al. \(2005\)](#), [Willer \(2013\)](#), [MacFarlane \(2011\)](#) and [MacFarlane \(2014\)](#). See [Huvenes \(2017\)](#) for more on disagreement as a state.

¹⁰See [von Fintel and Gillies \(2011\)](#) for a defense of contextualism. See [Lennertz \(2020\)](#) and [Yalcin \(2015\)](#) for responses. Also see [Plunkett and Sundell \(2013\)](#).

¹¹See [MacFarlane \(2011\)](#) and [Lennertz \(2019\)](#).

¹²See [MacFarlane \(2011\)](#), [Lennertz \(2019\)](#), and [Willer \(2013\)](#).

¹³The following formulation is equivalent to the one from [Lennertz \(2019\)](#) which is inspired by [Veltman \(1996\)](#).

A state supports a formula just in case updating the state with that formula results in the original state. This is to say that a state supports a formula just in case the information in that formula is already contained within the state.

Definition 3: Semantic Clauses

1. $s[p] = \{w \in s \mid w(p) = 1\}$
2. $s[\varphi \wedge \psi] = s[\varphi][\psi]$
3. $s[\varphi \vee \psi] = s[\varphi] \cup s[\psi]$
4. $s[\neg\varphi] = s - s[\varphi]$
5. $s[\text{Might } \varphi] = \{w \in s : s[\varphi] \neq \emptyset\}$

Update with an atomic sentence removes any worlds in the initial state where that sentence is not true. Conjunction is consecutive update, while disjunction is the union of the individual updates of the disjuncts. Negation removes the result of updating with the formula negated. Might φ runs a test on the information state, and returns the input state if update with φ doesn't crash. Otherwise, it crashes, (Veltman, 1996, 228). We can also define a Must operator as the dual of Might.¹⁴

Definition 4: Dynamic Consequence

$$\psi_1, \dots, \psi_n \Vdash_D \varphi \Leftrightarrow \forall s, s[\psi_1] \dots [\psi_n] \models \varphi$$

A sequence of formula dynamically entails another formula just in case every state is such that sequential update with those formula will support the other formula. This notion of consequence is often called *update-to-test consequence*.¹⁵ This definition of consequence is interdefinable with the following notion of consistency.

Definition 5: Dynamic Consistency

$$\psi_1, \dots, \psi_n \text{ is consistent}_D \Leftrightarrow \exists s, s[\psi_1] \dots [\psi_n] \neq \emptyset$$

Definition 6: Dynamic Inconsistency

$$\psi_1, \dots, \psi_n \text{ is inconsistent}_D \Leftrightarrow \psi_1, \dots, \psi_n \text{ is not consistent}_D$$

A sequence of formula is dynamically consistent just in case there is some state that can sustain the sequence of updates without resulting in the absurd state. Otherwise, it is dynamically inconsistent. Relevantly, dynamic consequence and consistency are order sensitive. For instance, the sequence $\text{Might } p, \neg p$ is

¹⁴Must $\varphi =_{\text{def}} \neg\text{Might } \neg\varphi$. On such an interpretation, Must is strong. While there are reasons to reject the strength of English “must” (see Karttunen (1972)), many working in the test semantics take “must” to be strong, so I shall follow suit. See von Fintel and Gillies (2007), Willer (2013), and Lennertz (2019) for such implementations. See von Fintel and Gillies (2007) and von Fintel and Gillies (2010) for arguments in favor of this position.

¹⁵Veltman (1996) calls this notion of entailment \Vdash_2 .

consistent_D while the sequence $\neg p$, Might p is inconsistent_D. However, no single state supports both $\neg p$ and Might p .

Even though semantic contents in a dynamic framework are not characterized in terms of propositions, propositions still do the representational heavy lifting, and many updates express propositions. The characteristic feature of updates like Might p is that they do not express propositions, and thus, we cannot even indirectly speak of their truth and falsity or their associated facts. Moreover, it is well known that the non-modal fragment of the semantics above behaves classically, but introduction of the Might operator introduces non-classical behavior. It will thus be useful to distinguish updates which express propositions from those that do not. Some content expresses a proposition when it has two properties: *distributivity* and *eliminativity*, (Groenendijk and Stokhof, 1991a, 57). All contents in the test semantics are eliminative, so some content is distributive if and only if it expresses a proposition.

Definition 7: Distributivity

φ is distributive \Leftrightarrow for all $s, s[\varphi] = \bigcup\{\{w\}[\varphi] \mid w \in s\}$

An update is distributive just in case updating a state with the formula will yield the same result as the union of the individual updates of the singleton set of each world in that state, (van Benthem, 1989, 364). All contents that do not include the Might operator are distributive, and hence, express propositions. Paradigmatically, contents of the form Might p are not distributive, and do not express propositions.¹⁶ Subsequent discussion will appeal to content that is propositional, which is to say, content that expresses update with a proposition.¹⁷ It will sometimes be useful to refer directly to a proposition expressed by a distributive update, rather than the update itself. We capture this as follows:

Definition 8: Proposition

For any φ that is distributive, $\llbracket \varphi \rrbracket =_{\text{def}} \{w \in W \mid \{w\}[\varphi] \neq \emptyset\}$

Distributive contents can always be associated with a set of possible worlds or *proposition*. The $\llbracket \cdot \rrbracket$ allows us to refer to this proposition directly. Crucially, $\llbracket \cdot \rrbracket$ is only defined for distributive contents.

Lastly, we need to define what it is for an agent to hold a belief. Here I provide Heim’s commonly accepted account of belief in dynamic frameworks.¹⁸ In traditional static frameworks, such contents will be propositions, but this won’t be so for dynamic theories. Instead, to have a belief is to have an information state that supports the update associated with the content of that belief.

Definition 9: Belief

¹⁶Notably, not all formulas of the form Might φ fail to be distributive. For instance, when the Might operator takes wide scope over a contradiction, e.g. $\text{Might}(p \wedge \neg p)$, the formula will be distributive. As Mandelkern (2020) observes, however, not every formula of the form $\varphi \wedge \neg \varphi$ is a contradiction within the test semantics.

¹⁷I will occasionally say of some propositional content that that content is true (or false). This is bit of shorthand to say that the content expresses a proposition that is true (or false).

¹⁸Instances of this can be found in Heim (1992), Yalcin (2011), Yalcin (2012b), Willer (2013) and Lennertz (2019) among others.

An agent a believes some content $\varphi \Leftrightarrow I_a[\varphi] = I_a$

An agent holds a belief just in case update with the contents of that belief will not change the agent’s information state. This is to say that the information contained in the update is already captured by the state, and thus, subsequent update has no effect. This can be equivalently defined in terms of support: an agent a believes some content φ if and only if $I_a \models \varphi$. Accordingly, an agent whose information state contains no p worlds will believe $\neg p$. Similarly, an agent whose information contains at least one p world believes *Might* p , and so forth. This account of belief is highly idealized insofar as agents who hold a belief will hold all the consequences of that belief. If an agent believes φ , and $\varphi \Vdash_D \psi$ then the agent also believes ψ . Similarly, if an agent is undecided with respect to the truth of some proposition $\llbracket \varphi \rrbracket$ then that agent’s information state will include both $\llbracket \varphi \rrbracket$ worlds and $\llbracket \neg \varphi \rrbracket$ worlds. Thus, if an agent is undecided, then that agent will believe *Might* φ and *Might* $\neg \varphi$.

The semantics above has been used to great effect to predict and explain various phenomena involving epistemic modals.¹⁹ Here, however, we isolate our attention to modal disagreements. Holmes believes $\neg m$ because updating Holmes’ information state, I_h , with $\neg m$ returns his initial state I_h ($I_h[\neg m] = I_h$). Watson believes that *Might* m and *Might* $\neg m$, since his information state I_w includes both m worlds and $\neg m$ worlds, and so update with either *Might* m or *Might* $\neg m$ (or both) will output I_w . However, should Holmes update his information state with Watson’s belief that *Might* m , Holmes’ state will crash. More formally, for any state s , if $s \models \neg m$ then $s[\text{Might } m] = \emptyset$. Holmes’ belief that Moriarty is not the murderer has ruled out all m worlds. Thus, updating with *Might* m will check for m worlds, but is doomed to fail, since his belief that $\neg m$ ensures that there aren’t any. Therefore, the content of Holmes’ belief is inconsistent _{D} with the content of Watson’s belief, and Holmes thereby disagrees with Watson.

If disagreement is characterized in terms of inconsistency, the test semantics is able to predict that Holmes disagrees with Watson. Cases of diachronic intrapersonal disagreement as well as eavesdropper cases can be explained in similar fashion. The issue, however, does not lie within the machinery of the test semantics. Rather, the alleged cases of disagreement that motivate the adoption of the test semantics are themselves motivated by the correctness judgments made by interlocutors like Holmes. This story is vindicated by the fact that Holmes judges the content of Watson’s belief to be incorrect. The test semantics, by itself, does not tell us anything about what it means for the content of a belief to be correct. Thus, we need to supplement the formal machinery above with a theory of correctness that explains why Holmes makes this judgment.

¹⁹Perhaps most notably, the test semantics predicts the infelicity of embedded epistemic contradictions. See von Fintel and Gillies (2007), Willer (2013), and Willer (2015) for details. See Yalcin (2007) for a thorough characterization of the problem of epistemic contradiction.

2 Correctness Conditions

2.1 Formative Correctness vs. Content Correctness

Proponents of the disagreement challenge to contextualism take the presence of modal disagreements to be evidence against contextualism. Such critics allege that within the Kratzerian contextualist framework, Watson’s belief that *Might m* expresses a second order description of the state he is in; something like, “For all I know, Moriarty might be the murderer.”²⁰ Thus, the content concerns Watson. If Watson is undecided as to whether Moriarty is the murderer, this content is true, and thereby correct. Moreover, this content is consistent with what Holmes believes. Unless Holmes misunderstands the nature of Watson’s beliefs, Holmes cannot plausibly disagree with respect to the content of Watson’s belief, nor can he believe that content to be incorrect. Since contextualism does not predict the content of Holmes’ belief and the content of Watson’s belief as jointly inconsistent, it is argued that contextualism fails to predict these cases of disagreement.²¹ Contextualists, however, may respond that purported disagreements like those between Holmes and Watson are not genuine disagreements, or at least, not the kind of disagreements that concern contents and inconsistency therebetween. To ameliorate these worries and motivate the objection, it is common to invoke the correctness judgments of interlocutors. For instance, Holmes can respond or believe that the content of Watson’s belief that Moriarty might be the murderer is wrong. The plausible assumption here is that when we judge the content of some belief or assertion to be incorrect, this is evidence of disagreement. The strategy of appealing to normative judgments about content strives to isolate the content of the belief or assertion as the locus of incorrectness, and thereby disagreement. This is to say that those who take the disagreement challenge to be an objection to Kratzerian contextualism insist that disagreement concerns content directly and place a demand on competing views that these disagreements manifest themselves in the semantics. This strategy places at least three constraints on proposals that can successfully answer the challenge. The first, and most well-recognized in the literature, applies pressure to commit to the claim that Holmes disagrees with Watson. The second is often less appreciated and holds that since the object of the normative judgement is the content of Watson’s epistemically modalized belief, it constrains the lines of response available to views which properly

²⁰Contextualists may not accept this characterization of their view. See, most famously, [Kratzer \(1977\)](#) and [Kratzer \(1981\)](#) for thorough explication of the position. See [DeRose \(1991\)](#) for a detailed account of contextualism concerning epistemic modals in particular.

²¹I oversimplify a bit here, as contextualists have argued that the contextually relevant state of information that the modal is sensitive to may include the information of agents other than the speaker/believer. Proponents of the disagreement challenge respond to this by proposing *eavesdropper cases*. Such cases are, allegedly, modal disagreements where the presence of the disagreeing agent is not known by the speaking/believing agent, and allegedly, no contextually supplied body of information is plausible. For examples of eavesdropper cases, see [MacFarlane \(2011\)](#) and [Khoo \(2015\)](#). The challenges I offer below do not require eavesdropper cases (though the solution I provide can explain them) so for simplicity, I avoid eavesdroppers in this discussion.

answer the challenge. In particular, it limits pragmatic and (arguably) metalinguistic explanations of the disagreement. Such restrictions are plausible, especially when we consider that disagreements tend to be *about* some content, and the disagreement in the present case is about whether or not Moriarty is the murderer. Thus, the correctness/incorrectness judgments of interlocutors like Holmes are a primary motivating factor for taking modal disagreements to be genuine disagreements, as well as insisting that they be explained by the contents of beliefs. Unlike the Kratzerian semantics, the test semantics does predict the content of Holmes' belief to be inconsistent with the content of Watson's. However, in order to utilize Holmes' correctness judgments as motivation, proponents of the test semantics will need to provide correctness conditions that explain why Holmes judges the *content* of Watson's belief to be incorrect.

This brings us to the third, and perhaps least appreciated, constraint, which concerns precisely what kind of normative judgment is being made. Lennertz (2019) makes a useful distinction in this regard, and I'll help myself to it here. Lennertz distinguishes between being correct/incorrect in believing some content and being correct/incorrect to believe some content. The first concerns whether the content of the belief is correct, while the second is an evaluation of how the belief was formed, (Lennertz, 2019, 4789).²² To see the difference, suppose that p is true. Then suppose that a believes that p , while b believes, falsely, that $\neg p$. Suppose further that a is doxastically reckless and has developed their true belief on radically insufficient evidence. Moreover, b is rational, has been doxastically responsible, and b 's best evidence supports $\neg p$. Thus, we can say that b is correct to believe $\neg p$, yet b is nonetheless incorrect in believing $\neg p$. The opposite holds for a , who is right in believing p but is wrong to believe p . While the process by which a formed the belief that p is unsatisfactory, and they are thereby wrong to believe p , the content of p is nevertheless correct, and a is correct in believing p .

Thus, there are at least two ways to normatively judge an agent with respect to their doxastic state. One concerns the process by which the agent formed their belief. We might call this *formative correctness*. An agent is formatively correct to believe some content iff their belief forming process satisfies the appropriate evidential and doxastic standards, whatever they may be. The other type of correctness, and the one relevant for disagreement, is determined wholly by the *content* of the belief. An agent is correct in believing some content iff that content is correct. The former notion of correctness concerns the inferential actions of the agent, and is sensitive to their evidence. The latter is not, and is entirely determined by the content of the belief. While both notions of correctness can also be applied to agents, formative correctness is sensitive to agential features over and above the content of the belief. On the other hand, whether an agent is correct in believing some content is determined solely by the content of the belief in question. As Lennertz argues, it is the second kind of correctness that is relevant for disagreement, (Lennertz, 2019, 4789). To see

²²Note that this distinction is not restricted to dynamic accounts and is equally applicable to static accounts. For instance, Wedgewood's thesis clearly concerns correctness in believing, and does not concern correctness to believe.

why, suppose that a third agent, c , knows that p . While c may respect b 's reasoning, she will disagree with b and will believe that the content of b 's belief is incorrect, since it is false. Alternatively, even if c does not respect a 's reasons for believing p is the case, c does not disagree with a about p . Regardless of how any of the agents formed their beliefs, c will agree with a and disagree with b . The quality of a and b 's evidence does not impact c 's capacity to agree or disagree with them about p .

Crucially, proponents of the disagreement challenge are committed to the claim that the relevant normative judgment made by Holmes concerns the content of Watson's belief, and does not concern Watson's evidence, or how he formed the belief. Holmes may very well judge that Watson is correct to believe *Might m* based on the evidence available to him. Nonetheless, he thinks Watson is wrong in believing *Might m* since that content cannot be correct when $\neg m$ is true. Proponents of the challenge that fail to appreciate this distinction will find that their challenge falls on deaf ears, as Holmes' evaluation of Watson's belief forming process places no pressure upon the semantics. Contextualist may respond that the disagreement, if any, concerns how Watson is responding to the evidence. This explanation can be provided without appeal to inconsistency of contents. Thus, any semantics which adequately responds to the challenge must be able to explain what it is for some content to be correct (and derivatively, what one is correct in believing), and should not concern formative correctness (what one is correct to believe.)

This sounds good when we limit our examples to propositional beliefs whose contents are truth-evaluable. True contents are correct, false contents are not, and therefore, agents are correct in believing true contents according to Wedgewood's principle. There is a lacuna in our story, however, should we insist that there are contents that are not truth-evaluable. The task at hand is to provide plausible conditions for content correctness within a dynamic framework that do not appeal exclusively to truth. In addition, these conditions should also explain why Holmes judges the content of Watson's belief to be incorrect.

2.2 Candidates for Content Correctness

The literature features a few different accounts of a correct belief that are compatible with the test semantics. While some dynamic accounts appeal to the distinction between content correctness and formative correctness, there are not, to my knowledge, any dynamic accounts that explicitly characterize the conditions for these different kinds of correctness.²³ As we shall see, this leads to difficulties. The proceeding section explores extant proposals, and considers whether they are, in fact, accounts of content correctness or formative correctness. After considering each, I explain why each account is, unfortunately, unsatisfactory.

²³Lennertz (2019), for instance, appeals to the difference between being wrong to believe and being wrong in believing but does not provide an account of precisely what it is to be wrong in believing within the test framework.

2.2.1 Evidential Correctness

Willer (2013) deploys the test semantics to engage the puzzle of modal disagreement. Willer’s strategy involves adding additional complexity to the semantics described above in supervaluationist fashion in order to capture how epistemic modals express attention towards a possibility. The result is a sophisticated development of the test semantics that addresses a battery of problems involving epistemic modals. After proposing the semantics, Willer addresses an objection against non truth-conditional accounts of epistemic modals from MacFarlane (2011). MacFarlane argues that such accounts are unable to explain the incorrectness judgments which motivate disagreements and retractions, (MacFarlane, 2011, 158). Willer recognizes the importance of correctness judgements, but rejects the claim that no account of correctness is available for the test semantics. He does this by providing his own. Call this condition *evidential correctness*.

Evidential Correctness: A belief is correct if and only if it would not commit someone with the believing agent’s evidence to factual error, (Willer, 2013, 64).

According to evidential correctness, an agent is correct to believe some content just in case no one equipped with the same evidence as the believing agent would be committed to factual error. Willer also provides a definition which explains what “factual error” amounts to.

Definition 10: Truthfulness

An information state s is truthful $\Leftrightarrow @ \in s$

An information state is truthful just in case the state has not ruled out the actual world, (Willer, 2013, 56). The actual world is the privileged possible world where all of the propositions that are true at that world, are, in fact, true. An information state that contains any information which is false will rule out the actual world. Thus, an agent is committed to factual error just in case their information state is not truthful.

Evidential correctness is compelling and appropriately broad. It is somewhat intuitive that a “might” belief is correct when the prejacent is compatible with the evidence. Thus, in many cases, it does not attribute error to those who refrain from holding a fully-fledged propositional belief when the evidence does not support that belief. Further, when an agent’s evidence does support some belief, φ , and the agent fails to make the appropriate inference and still believes *Might* $\neg\varphi$, evidential correctness rightly predicts that the believing agent will be incorrect. This is because a more discerning agent with the same evidence would have concluded φ , and thus would be committed to factual error if they believed *Might* $\neg\varphi$.²⁴ In addition, evidential correctness applies to propositional and non-propositional beliefs alike. If the proposition associated with a belief is false, then it directly leads to factual error. In addition, certain epistemic modal beliefs that do not express propositions, e.g., “It must be that p ” can lead to factual error and can be incorrect as a result. Lastly, while evidential correctness does not appeal directly to truth, the truth remains relevant in that the agent’s

²⁴The only state that supports both φ and *Might* $\neg\varphi$ is \emptyset , which is not truthful.

beliefs should not conflict with the facts. Thus, evidential correctness remains suitably grounded in reality.

Prima facie, evidential correctness seems intuitive, and attractive. It also appears to make the right predictions in the case involving Holmes and Watson. Suppose that Holmes and Watson have just jointly investigated the same crime scene and have the same evidence. Holmes, in classic fashion, has made inferences that Watson has missed, despite having access to the very same evidence. Based on this evidence, Holmes has miraculously deduced that Moriarty is not the murderer. Were he to maintain these beliefs while also believing the content of Watson's belief, he would be committed to factual error, since $\neg m$ is inconsistent_D with Might m . Thus, the counterfactual conditions for evidential correctness do not hold for Watson's belief, since there is someone, namely Holmes, who has Watson's evidence, yet would be committed to factual error if he held Watson's belief. If Holmes adopted Watson's belief, he would be in the absurd state, guaranteeing that the actual world was removed from his state, and thereby committing him to factual error. Watson is thereby incorrect. Equally importantly, he is incorrect in a way that is accessible to Holmes, allowing Holmes to judge that that Watson is incorrect.

Despite the auspicious start, cracks appear under closer scrutiny. There are two substantive problems with evidential correctness that prevent it from playing the dialectical role that Willer intends it to play. The first is that while evidential correctness makes the right predictions when Holmes shares the same evidence as Watson, evidential correctness fails to make the requisite predictions outside of such carefully engineered scenarios. In particular, it fails in cases where agents do not share the same evidence. For instance, suppose that the disagreement between Holmes and Watson takes place after the two have each independently investigated separate sections of the crime scene. Holmes' evidence supports $\neg m$ while Watson's provides no evidence for or against m . Each are unaware of the other's evidence and do not share the same evidence. According to evidential correctness, Watson's belief is incorrect just in case someone with Watson's evidence would be committed to factual error, and this is supposed to explain why Holmes can utter, "No, you're wrong," if Watson asserts Might m . The problem here is that Holmes is entirely unaware of Watson's evidence, and thus, if his normative judgment about Watson's belief is guided by evidential correctness, he will be unable to judge whether someone with Watson's evidence would be committed to factual error. Notably, such scenarios still allow for modal disagreements, which require normative judgments, and it strikes as perfectly felicitous for Holmes to utter, "No, you're wrong." This means that Holmes' assessment that Watson is incorrect is coming from some other principle, and is not explained by evidential correctness.

The problem is not merely one of epistemic access. Suppose further that Holmes can read minds, and is aware of Watson's evidence. Even if Holmes can read Watson's mind, if his judgments are guided by evidential correctness, Holmes should conclude that Watson's belief that Might m is correct. Since Watson's evidence is silent with respect to m , no one with Watson's evidence would be committed to factual error. Nonetheless, Holmes can still felicitously

and justifiably utter, “No, you’re wrong.” Evidential correctness fails to predict this judgment. Thus, evidential correctness does not make the requisite predictions with respect to the normative judgements required for the disagreement. If Holmes believes $\neg m$, he will judge any Might m belief to be incorrect, whether he is aware of the evidence or not.

Recall further that disagreement concerning states of information does not require that conversation manifest, and the absence of conversation does not prevent agents from disagreeing. Thus, Holmes need not even speak with Watson, nor have any ideas about Watson’s evidence, in order to disagree. Just like b and c above, Holmes disagrees with Watson because he judges some content to be incorrect, and he is able to do this regardless of the evidence that supports that content. Evidential correctness fails to explain this, and thus fails to motivate the disagreement challenge in cases where disagreeing agents do not share evidence. Things are made worse by the fact that scenarios where disagreeing agents do not have access to one another’s evidence are common. In fact, we often engage in active information exchange precisely because we do not share the same evidence. One may even suspect that disagreement between agents that do not share evidence are more common than ones where they do. At the very least, cases where evidential correctness fails aren’t *recherché*, nor need they involve complexities like eavesdropping. Such cases are common if not the norm.

The problem becomes more pronounced when we consider why MacFarlane demands an account of correctness/incorrectness in the first place. He insists that one be provided precisely because these kinds of normative judgments are what motivate the disagreement challenge. If the account of correctness provided is unable to supply the judgments that motivate the challenge, then MacFarlane, and proponents of the challenge more generally, should remain unsatisfied. A satisfactory account of correctness will be able to predict and explain, in all instances of modal disagreement, that the disagreeing agent is able to judge their interlocutor to be incorrect. Evidential correctness does not meet this criteria.

The above is suggestive of the second, deeper problem. Under Willer’s account, some belief is evidentially correct based on some body of evidence, and what inferences would be made by other agents in possession of that evidence. This sensitivity to evidence suggests that evidential correctness is not a condition on the correctness of content, but is rather a condition for formative correctness. As we saw above, formative accounts of correctness do not motivate the disagreement challenge. Importantly, while Willer discusses correctness, he does not distinguish content correctness from formative correctness, and I suspect that there is a conflation of the two. These suspicions are supported by the following from Willer.

“Correctness as characterized above turns on an individual’s evidential situation and thus we leave room for the possibility that, given adequate variation in what is known, [Watson] correctly believes that [Moriarty might be the murderer], while [Holmes] correctly be-

believes that [Moriarty is not the murderer]... But we are well advised to bear in mind that correctness thus understood cannot serve as a reliable guide to the semantics and pragmatics of epistemic modals,” (Willer, 2013, 64-65).²⁵

The above makes explicit the fact that, under evidential correctness, whether an agent is correct to believe something is determined by the evidence available to them. This is, arguably, *the* characteristic feature of formative correctness. Moreover, just as Willer claims, as one’s evidence changes, so changes what one is correct to believe. This explains the differing judgments of correctness between Watson and Holmes, but is also characteristic of an account of formative correctness. However, as Lennertz convincingly argues, the normative notion relevant to the disagreement challenge is not formative correctness, but content correctness. The disagreement challenge does not concern what Holmes thinks Watson is correct to believe, but what he thinks he is correct in believing.

I take the final sentence of the quote above to mean that, since evidential correctness concerns which contents are correct to believe given certain bodies of evidence, and since the same belief can be correct relative to some body of evidence, and incorrect relative to another, that judgments of correctness should play no deep role in our semantic account of epistemic modals. Later in the paper, Willer precisifies the point.

“The point is that such correctness criteria do not serve as a reliable guide to the semantics and pragmatics of epistemic modals. *A speaker’s judgment that so-and-so might be the case may be correct—the speaker may be responding correctly to his or her evidence—and still be rejected by a better informed assessor... A hearer’s reaction to a judgment involving might or must is then expected not to be guided by considerations of correctness but rather by the judgment’s update effect on the hearer’s information state. We have seen that this approach easily accounts for the observation that ordinary speakers assess judgments involving might by testing them against their own perspective...*” (Willer, 2013, 87) (emphasis added).

Willer argues that correctness judgments are not what motivates Holmes’ rejection of Watson’s belief. Instead, Holmes seems to judge Watson’s belief based on his own evidence. Watson’s belief is incompatible with Holmes’ information, and thus Holmes rejects Watson’s proposal. This explains why Holmes rejects Watson’s belief but it fails to explain why Holmes thinks Watson is incorrect in believing this. Willer thinks this isn’t a problem since Holmes’ ability to reject Watson’s claim is not based on Holmes’ assessment of its correctness (Holmes might think Watson is perfectly correct to believe *Might m*), but instead rejects Watson’s claim based on the fact that it is incompatible with Holmes’ information. Unfortunately, I do not think that this is a move Willer is able

²⁵I alter Willer’s quote by exchanging Holmes and Watson with Willer’s example subjects, Alex and Mary.

to make. Recall that the disagreement challenge relies fundamentally on the argument that Holmes' normative judgment of the content of Watson's belief is what ensures that he disagrees with Watson about that content, and is not merely rejecting his claim.²⁶ Should we trivialize the role of Holmes' normative judgment in modal disagreements, as Willer suggests, then the contextualist is free to trivialize it as well. They may explain Holmes' rejection pragmatically or for other non-semantic reasons, and are under no pressure to modify the semantics. Thus, should Willer pursue this strategy, his claim that dynamic semantics can answer the challenge doesn't amount to much, since the presence of the challenge offers no pressure to abandon contextualism in favor of a dynamic alternative. This is to say that, by my lights, Willer's proposal does not answer MacFarlane's complaint. Willer provides an account of correctness only to concede that this account of correctness plays no role in Holmes' rejection of Watson's claim. This, however, is precisely why MacFarlane demands an account of correctness in the first place.

Willer's assessment that evidential correctness cannot be what motivates Holmes' judgment seems to dovetail with my observations that it fails to make the requisite predictions. This explains why Willer thinks that correctness judgments should not inform our semantic theorizing. Unfortunately, the disagreement challenge insists precisely that we do allow correctness judgments to inform our theorizing. Willer's comments, most notably the italicized portion of the quote above, suggest, however, that Willer has only considered formative notions of correctness which are sensitive to evidence. As Lennertz observes, however, the disagreement challenge concerns correctness of content. It should thus be unsurprising that evidential correctness fails to motivate Holmes' judgment. It is the wrong kind of correctness. My suspicion is that Willer has attempted to recruit a formative norm, when only a content norm will suffice, and as a result, has undermined the challenge he has striven to solve.²⁷ This also points us in a more fruitful direction. What we need is an account of correctness for non-truth-conditional content that does inform our semantic theorizing. This account should also vindicate the claim that Holmes thinks Watson is incorrect in believing *Might m*, without appeal to evidence or any other formative norms. Such an account of correctness will allow the dynamic semanticist to answer the challenge in a way that the contextualist cannot and thereby motivate the view. We should thus explore options that concern content alone.²⁸ Two options readily suggest themselves.

²⁶We can often reject things that we do not judge to be incorrect.

²⁷This should also make clear that potential variations of evidential correctness will similarly fail to answer the challenge. So long as such proposals are sensitive to evidence, they will be formative accounts of correctness, rather than accounts of the correctness of content.

²⁸This argument does not amount to a wholesale rejection of evidential correctness. It may well be a perfectly appropriate formative norm. My contention is that evidential correctness cannot be the kind of correctness operative in modal disagreements.

2.2.2 Weak Correctness

With respect to content correctness, an agent’s ability to be correct in believing some content depends entirely on the semantic value of that content. Not so for formative correctness, where an agent can be perfectly formatively correct, even if the content is not. This suggests that conditions for content correctness should concern content directly and should not appeal to evidence or inferential mechanisms. This simultaneously explains why evidential correctness failed, while pointing us in a more productive direction. I next consider a proposal in this spirit, which I call *weak correctness*.

Weak Correctness: If an agent a holds a belief with content φ the content of that belief is correct if and only if $@ \in I_a[\varphi]$.²⁹

Weak correctness holds that a belief is correct if and only if that belief does not commit the believing agent to factual error. This is to say that a belief is correct just in case that belief does not cause the believing agent to believe falsely.

Weak correctness has some immediate advantages. The first is that it does not appeal to evidence or inference, ensuring that it characterizes correctness of content. The second is that, like evidential correctness, it applies to propositional and non-propositional beliefs alike. If the content of some proposition is true, then it will be correct. If it is false, it will be incorrect. Things change with respect to epistemic modal beliefs, which are neither true nor false. Consider the content *Might p* . A belief in this content will be correct if the believing agent’s information state contains at least one p world. It will be incorrect if it does not contain any p worlds. More generally, we see that some epistemically modalized content is correct just in case the test performed by the modal is passed. This both explains why one is incorrect in believing *Might p* when one believes $\neg p$, but it also explains the intuition that “might” beliefs can be correct when an agent is undecided as to the truth of some proposition.

The obvious problem, however, is that weak correctness universally fails to vindicate Holmes’ judgment that Watson’s belief is incorrect. Watson’s belief that Moriarty might be the murderer does not commit him to factual error, and is thereby correct. Moreover, even though Holmes believes $\neg m$, Holmes will recognize that Watson’s belief that *Might m* will not commit him to factual error. Thus, Holmes should judge Watson to be correct in believing *Might m* . This problem generalizes to all instances of modal disagreement and universally fails to predict the incorrectness judgments that motivate the problem. Thus, whatever its advantages, weak correctness cannot be the correctness condition that motivates the disagreement challenge.

2.2.3 Strong Correctness

Willer (2013) mentions an alternative in a footnote on pp.63, which appeals to a proposal from Yalcin (2011). Yalcin proffers a condition for correct assertion

²⁹Here and henceforth, I characterize contents in terms of formulas as opposed to updates, but this choice is purely practical, and nothing hangs on it. Definitions concerning content can be easily altered to instead appeal to updates rather than formula.

called *advisability*, but we can instead articulate this as a correctness condition for the content of belief and explore the results. Willer does not pursue this condition in depth, but, in light of my objections, we can take a closer look.

Strong Correctness: A belief is correct if and only if an agent who knows all of the relevant facts holds this belief, (Willer, 2013, 63).³⁰

According to strong correctness, a belief is correct just in case some idealized agent who already knows the relevant facts holds this belief. Thus, if Holmes and Watson are debating about m , the idealized knower will know whether or not m is the case. Strong correctness is significantly stingier than weak correctness when it comes to whose beliefs are correct, specifically with respect to “might” beliefs. Just like weak correctness, any commitment to factual error will result in a belief that is incorrect. However, weak correctness was permissive with respect to some epistemic modal beliefs. Even if some proposition p was false, weak correctness allowed for beliefs like *Might* p to be correct, so long as they did not commit the believer to factual error. Not so for strong correctness. According to strong correctness, if p is false, then an agent who knows the relevant facts will not believe *Might* p , and so, any such belief is incorrect. Thus, it is not enough to merely fail to believe falsely. In order to satisfy strong correctness, one must believe truly.

We can now apply strong correctness to the disagreement between Holmes and Watson. Holmes believes $\neg m$, which is to say that Holmes believes the facts are such that an idealized agent would believe $\neg m$, and would not believe *Might* m . This explains why Holmes judges the content of Watson’s belief to be incorrect. Note that this applies even if it turns out that Holmes himself is mistaken, and Moriarty is in fact the murderer. In virtue of believing $\neg m$, Holmes believes that the facts are such that $\neg m$ and subsequently believes that an idealized agent would also believe $\neg m$.³¹ This is looking good, since it explains Holmes’ judgment that *Might* m is incorrect. Consequently, this predicts that Holmes will judge any agent who believes *Might* m to be incorrect in believing this content. One can change the scenario as freely as one likes, adding differing evidence, eavesdroppers, or whatever else. Strong correctness predicts that Holmes will judge any belief that *Might* m to be incorrect, as is required to motivate the puzzle.

Unlike weak correctness, strong correctness predicts that Holmes judges the content of Watson’s belief to be incorrect, maintaining the initial motivation for modal disagreement. Unlike evidential correctness, strong correctness concerns content exclusively, and is thus an appropriate condition on correctness of con-

³⁰This definition appeals to a locally idealized knower with access to the relevant facts. Those with concerns about the work that relevance plays may appeal to an absolutely idealized knower who knows all of the facts. What follows is equally applicable to both characterizations, so the distinction is not emphasized.

³¹Put another way, when it comes to our propositional beliefs, we tend to think of ourselves as locally idealized agents, since we think our own propositional beliefs to be true. We can be mistaken, however, about what an idealized agent would believe. This also aligns nicely with the observation made by Willer and others that agents tend to test “might” beliefs “against their own perspective,” (Willer, 2013, 87).

tent. In addition, it allows Holmes to make the requisite correctness judgment about the content of Watson’s belief in all cases, regardless of evidence. Strong correctness thus avoids the challenges faced by evidential and weak correctness and appears to be the most attractive option thus far. Unfortunately, it also makes a multitude of less attractive predictions concerning the beliefs of individuals who are undecided. Recall that for the test semantics, to be undecided about some proposition p involves, among other things, having an information state that can sustain update with p as well as $\neg p$. Unlike traditional propositional accounts, where indecision is the absence of belief, our current framework ensures that indecision requires at least some beliefs. If an agent is undecided with respect to p , then their information state will support $\text{Might } p \wedge \text{Might } \neg p$, and thus, the agent will believe $\text{Might } p \wedge \text{Might } \neg p$.³² It doesn’t matter whether this belief is occurrent, or whether the agent recognizes that they have it.

In a purely propositional framework, strong correctness appears plausible and attractive.³³ However, in frameworks where indecision entails belief in some content, there are some rather extreme consequences.

Ubiquity of Error: For any propositional content φ , if $@ \in \llbracket \varphi \rrbracket$, then any doxastic agent who fails to believe φ , and thereby believes $\text{Might } \neg\varphi$ believes some content that is incorrect and thus, is wrong in believing $\text{Might } \neg\varphi$.

According to strong correctness, any agent who fails to believe truly with respect to any proposition is guaranteed to hold a belief that is incorrect. This includes both agents who believe falsely, as well as agents who are undecided. An immediate consequence is that every actual doxastic agent is incorrect in holding at least some of their beliefs. Both weak and evidential correctness missed crucial cases of incorrectness and thus failed to motivate the disagreement challenge. Strong correctness avoids this problem by painting with an especially broad brush and characterizing any content that is not consistent with the facts as incorrect. It doesn’t matter whether the inconsistency is the product of indecision or false belief.

It is important to appreciate that the traditional, static framework operates differently. According to Wedgewood’s truth norm, the content of a belief is correct just in case it is true. False beliefs are incorrect, but the absence of belief involves no content that can be considered correct or incorrect by the truth norm. Classically, to be undecided with respect to some proposition is to not have a belief towards that proposition, and indecision requires no belief at all. While indecision may be correct in some circumstances and not others, this correctness is not governed by any content norm—there is no content—and is instead governed by some formative norm of belief. As a result, the classical framework does not predict such large-scale error in content. Unlike classical frameworks, the test semantics treats undecided agents as believing some content. When coupled with strong correctness, the test semantics is able to explain how the

³²The same agent will also believe $\text{Might } p$ and $\text{Might } \neg p$.

³³If we interpret strong correctness as suggested in footnote 27, then it is equivalent to Wedgewood’s truth norm in frameworks that are propositional.

content of certain “might” beliefs are wrong. The consequence, it turns out, is that a whole lot of content, including all indecisive content, is wrong.³⁴

This result is not as extreme as it may appear at first blush. False content and indecisive content each fail to accurately represent the world and are thereby incorrect. More importantly, however, strong correctness makes no judgments about what one is correct to believe. Strong correctness only governs what one is correct in believing. Watson is very plausibly correct to believe *Might m*. Strong correctness takes no stand on this issue and merely predicts that he is incorrect in believing it. The tension between the pressure to respond appropriately to evidence (formative correctness) and the pressure to believe truly (content correctness) is nothing new, nor is it unique to dynamic frameworks. Accordingly, the ubiquity of error may well be a bullet worth biting. Unfortunately, when it is coupled with another consequence of strong correctness, the result is much more difficult to embrace.

Homogeneity of Error: There is no normative difference between indecisive content and false content. Both are incorrect in exactly the same way, for exactly the same reason.

The homogeneity of error arises from the fact that strong correctness does not discriminate between different kinds or degrees of content correctness. Since being correct in believing some content is wholly determined by the correctness of that content, it follows that believing indecisive content is just as incorrect, and is incorrect in precisely the same way, as believing falsely. Thus, there is no difference in correctness between believing indecisive content and believing false content.

The test semantics recognizes contents that express indecision directly, without describing the agent as being in an undecided state. Accordingly, correctness conditions for such contents seem obliged to distinguish between the manner in which one is incorrect in believing indecisive content, and the manner in which one is incorrect in believing falsely. Strong correctness fails to do this. What’s more is that this distinction is clearly suggested by the test semantics itself. To see this, we can enrich the scenario with Holmes and Watson. Suppose that the interaction also includes a third agent, Lestrade, who believes that Moriarty is the murderer (*m*). Suppose further, for objectivity, that $@(m) = 0$. Observe that while both Watson and Lestrade are, in some sense, incorrect, they are in radically different doxastic situations. Watson, in virtue of being undecided, can still consistently come to know the truth that Moriarty is not the murderer. Watson believes *Might m* \wedge *Might* $\neg m$, however, the sequence, $\langle \text{Might } m \wedge \text{Might } \neg m, \neg m \rangle$ is dynamically consistent.³⁵ Thus, the contents of Watson’s beliefs do not impede him from arriving at the truth. Lestrade’s situation is different. Lestrade believes *m*, which is dynamically inconsistent

³⁴Some content φ is indecisive with respect to some proposition $[\psi]$ \Leftrightarrow the only states that can sustain update with that content are states that include at least one world w s.t. $w \in [\psi]$ and at least one world w' s.t. $w' \in [\neg\psi]$. Thus, *Might p* \wedge *Might* $\neg p$ is indecisive with respect to $[[p]]$.

³⁵This is because many *Might* updates are not *persistent*.

with the factual update expressed by $\neg m$. Thus, Lestrade’s false belief not only ensures that his information state will not be truthful, but it actively prevents him from productively updating with the truth. Thus, Lestrade will need to reject his belief that m if he is to have any hope of arriving at the truth. It should be clear that, based on the semantics alone, The content of Lestrade’s belief is incorrect in a way that the content of Watson’s is not. The test semantics and, hopefully, our intuitions suggest a normative distinction between indecisive content and false content. Strong correctness, the content governing norm responsible for making such individuations, is not sensitive to this distinction, and treats indecisive content as wrong in the same way as content that is false.³⁶ All doxastic sins are weighed in equal measure.

One may again be tempted to take the homogeneity of error in stride. This is made difficult when we consider that the homogeneity of error coupled with the ubiquity of error entails that every doxastic agent that is undecided is not only wrong in holding at least some of their beliefs but wrong in the same way as those who believe falsely.

The final problem arises when we consider how normative judgments motivate our actions. Upon realization that some content is incorrect, agents should cease to believe that content, and revise their beliefs accordingly. We also expect this from our interlocutors. Holmes does not make a normative judgment about Watson’s belief with the hope that Watson will maintain it. Rather, Holmes is trying to convince Watson to believe $\neg m$. If indecisive content is just as wrong as false content, and we recognize that we believe some indecisive content, then we should judge it to be wrong. We should thus react to our own indecisive beliefs in the same way that we react when we realize that our beliefs are false, by rejecting them.

Forced Decision: An agent who judges the content of an indecisive belief to be incorrect should reject the content of that belief.

To be undecided about p requires that an agent believe $\text{Might } p \wedge \text{Might } \neg p$. To be consciously undecided is to be aware that *inter alia* one believes $\text{Might } p \wedge \text{Might } \neg p$. However, despite the fact that an agent doesn’t know whether p or $\neg p$ they will know that one is true and the other is false. As such, $\text{Might } p \wedge \text{Might } \neg p$, is incorrect according to strong correctness. Moreover, the homogeneity of error ensures that the incorrectness of this content is no different than the incorrectness of false content.

This turns out to matter a great deal when we consider what happens when we know that the content of a belief is incorrect. Consider a straightforward case of false belief. Suppose a believes that p . Later, a discovers that p is false. They thus judge their belief that p to be incorrect. In response to this, they revise their beliefs to state such that $s_a \models \neg p$. So far so good, but strong correctness treats indecisive contents in exactly the same way. If a consciously believes $\text{Might } p \wedge \text{Might } \neg p$, then they should judge it as incorrect, and thereby revise their

³⁶Data gathered in Khoo (2015) suggests that speakers are sensitive to this distinction, or one very much like it, insofar as experimental subjects are inclined to judge “might” beliefs with false prejacent as wrong, but not false.

beliefs such that $s_a \models \neg(\text{Might } p \wedge \text{Might } \neg p)$. But the only states that support $\neg(\text{Might } p \wedge \text{Might } \neg p)$ are states that support p and states that support $\neg p$. The only state supporting both is the absurd state, \emptyset . Thus, according to strong correctness, *a ought to guess*. This is neither what we should want, or expect from an epistemic norm governing content. What is especially interesting is that, as far as strong correctness is concerned, there is no risk in this guess. Any state that supports $(\text{Might } p \wedge \text{Might } \neg p)$ is guaranteed to be incorrect. However, by guessing, the agent is only increasing their chances of being correct. If they guess wrong, they are no worse off than if they stayed undecided. If they guess right, they have disabused themselves of believing incorrectly.³⁷ Accordingly, should we adopt strong correctness for the test semantics, we would expect undecided agents to be pressured to guess in an attempt to avoid error. This is neither what we observe in everyday thought and talk nor what we should want from rational agents.³⁸

By my lights, things look fairly dire for strong correctness. The idea that every undecided agent holds an incorrect belief may be surprising but is a manageable pill to swallow. However, the idea that such beliefs are just as wrong as false beliefs, coupled with the further claim that undecided agents should guess in order to avoid incorrectness is simply too much to choke down. In its current form, strong correctness appears untenable.

2.2.4 Where to Go from here

This section has argued that evidential correctness, weak correctness, and strong correctness suffered from problems. Evidential correctness was a formative norm, rather than a content norm. In addition, both evidential correctness and weak correctness failed to yield the normative judgments that motivate the disagreement challenge. Strong correctness succeeded in both these regards, but yielded doxastically implausible consequences. The problems with strong correctness seem to stem from the fact that indecision involves belief within the test semantics. This suggests one of two strategies to avoid the problem.

³⁷We might be inclined to separate two normative notions for doxastic action: one that drives us to believe correctly, and one that deters us from believing incorrectly, as recommended in James (1896). We may further place these norms in a hierarchy where one trumps the other. For example, If we prefer correct beliefs more strongly than we disprefer incorrect ones, one would be disposed to guess, perhaps in runaway fashion, because this would yield a net positive. To avoid rampant guessing, we may instead choose to flip the order of the hierarchy, and make incorrect beliefs more undesirable than correct beliefs are desirable. One would expect this to prevent rampant guessing, since the negative cost of getting it wrong outweighs the positive cost of getting it right. In a propositional framework, this is exactly what you'll get, making the more conservative hierarchy of the preferences the better option. This fails within the current framework, however, since it is the pressure to avoid incorrectness which motivates guessing. Thus, within the current framework *either ordering of the norms will yield a guess*, so hierarchies like these will not avoid the puzzle.

³⁸Importantly, strong correctness is a content norm, and other prudential and formative norms may ultimately prevent a guess, all things considered. However, pressure to guess simply should never arise from a norm that concerns content. While we should expect tensions between various different epistemic norms, norms governing content should not have this kind of internal tension.

The first is to reject the claim that believing $\text{Might } p \wedge \text{Might } \neg p$ is a necessary condition for indecision about p . This option is explored, with different motivations, in [Willer \(2013\)](#) and [Yalcin \(2018\)](#). While each of these options will mitigate certain instances of the problems, they continue to persist. I address each of these strategies briefly in the next section. The second, more attractive, option is to recognize that indecisive content is wrong in a different way than false content, and alter our account of correctness accordingly. This option is developed in §4.

3 Rethinking Indecision

In an effort to make my argument as general as possible, I have provided a fairly simple version of the test semantics. The literature, however, features more sophisticated proposals, like [Willer \(2013\)](#) and [Yalcin \(2018\)](#), that allow for cases where one is undecided about p , but need not believe $\text{Might } p \wedge \text{Might } \neg p$. Since my criticisms of strong correctness assume this connection, these accounts will have to be addressed.

The version of the test semantics in [Willer \(2013\)](#) is complex and resists brief summary. A crucial feature, however, is that information states have added structure that alters the update effects of epistemic modals. For Willer, information states are not sets of worlds, but sets of sets of worlds. Call each set of worlds within a state a *substate*. When an update is performed on a state, updates are applied to each substate, and then substates are aggregated. The result is that epistemic modals do not simply test for compatibility with the prejacent. Rather, a state will only support $\text{Might } p$ on condition that every substate contains at least one p world. This is intended to capture the idea that sentences like $\text{Might } p$ express more than mere compatibility with p , but that p is taken seriously, or matters, in inquiry, ([Willer, 2013](#), 56). Accordingly, an agent’s information state will only support $\text{Might } p$ if that agent is disposed to take p seriously as a possibility.

A promising result for Willer’s semantics is that agents who are not disposed to take both p and $\neg p$ seriously will not believe $\text{Might } p \wedge \text{Might } \neg p$, and can thereby avoid error. This can include agents who are undecided about p as well as agents who simply never have considered p .³⁹ The result is that a large portion of undecided agents will no longer be incorrect in believing $\text{Might } p \wedge \text{Might } \neg p$, since they do not need to believe that in order to be undecided. We see a similar result with forced decision. Since not all undecided agents will believe $\text{Might } p \wedge \text{Might } \neg p$, they will not be judged as incorrect, nor will they be pressured to guess.

This interaction with strong correctness is undeniably an improvement over what was previously discussed. However, strong correctness still yields unhappy

³⁹The latter applies to many, but perhaps not all agents who have not considered p . For Willer, considering something to be a relevant alternative is a matter of disposition, ([Willer, 2013](#), 50). This, in principle, leaves open the possibility for an agent who has never considered p to be disposed to take p seriously, upon consideration.

results when coupled with Willer’s proposal. In particular, the account falters in instances where undecided agents take both alternatives seriously. Unfortunately, many of the most important and difficult instances of deliberation are precisely this way. Watson, for example, is aware of Moriarty’s past crimes, and takes the possibility that Moriarty is the murderer seriously. However, he is not so foolish as to believe that this is the only possibility worthy of attention and takes the possibility that someone else is the murderer (and Moriarty is not) equally seriously. In such cases, the negative effects of the homogeneity of error remain, as the content of Watson’s belief that $\text{Might } m \wedge \text{Might } \neg m$ will be incorrect. Forced decision also looms, although with a new wrinkle. In the simple test semantics, the only way to reject the belief that $\text{Might } p \wedge \text{Might } \neg p$ was to decide. Willer’s account offers a second way to reject the belief: cease to take at least one of the possibilities seriously. The good news is that agents who believe $\text{Might } p \wedge \text{Might } \neg p$ are no longer forced to guess. The bad news is that the new option isn’t much better, as ceasing to take a relevant alternative seriously sounds similarly irrational. One might try to bite the bullet here, and think that this is still a better option. In doing so, one might embrace a weaker principle whereby agents should not consider incompatible alternatives seriously at the same time. If this principle holds, we would expect sentences like the following to be infelicitous:

- (1) It might be raining, and it might not be.

However, we felicitously utter sentences like (1) regularly, and the principle strikes as untenable. Thus, the proposal made in Willer (2013) still suffers from critical problems associated with strong correctness.

Yalcin (2018) does not propose a semantics, but rather, an account of belief where belief is *question sensitive*. On this account, beliefs are relativized to questions or *resolutions* of logical space, (Yalcin, 2018, 30). The idea is that agents can fail to recognize certain distinctions in the logical space. When an agent does recognize a distinction, her logical space will represent at a higher resolution, individuating possibilities that it previously did not. The account is intended to solve various puzzles concerning logical omniscience, but the view allows for an agent to be undecided about p , without believing $\text{Might } p \wedge \text{Might } \neg p$. For instance, if an agent has never considered the issue of whether or not p , their logical space will not individuate p worlds from $\neg p$ worlds. The same goes for agents that lack some relevant concept crucial to understanding p , (Yalcin, 2018, 34). The pertinent result, with respect to strong correctness, is that agents whose information state lacks the appropriate resolution for some proposition p , will not believe $\text{Might } p \wedge \text{Might } \neg p$, despite being undecided. The result is much the same as it was for Willer, in that the ubiquity and homogeneity of error are substantively mitigated. Undecided agents who have never considered whether p will not believe $\text{Might } p \wedge \text{Might } \neg p$, and thereby will not hold an incorrect belief.

Unfortunately, this strategy is of no help for agents whose logical space does respect the appropriate resolution. Assuming that agents know what they are

saying, this will include any agent who asserts a “might” sentence. Thus, agents who ask whether p , but remain undecided, will believe $\text{Might } p \wedge \text{Might } \neg p$ and run into all of the same problems brought on by strong correctness. There is a similar wrinkle associated with forced decision, where an undecided agent is not forced to decide, and instead, may choose to cease asking the question, and coarsen the resolution of their logical space in the face of the incorrectness of their indecisive belief. This result is decidedly unattractive.

Thus, we observe that separating indecision about p from believing $\text{Might } p \wedge \text{Might } \neg p$ only helps to a limited extent, and does little to solve the hard cases. Relevantly, neither solution appears to help us with Watson and Holmes. It is worth emphasizing that nothing I’ve said here amounts to rejection or criticism of either view, and there are strong independent reasons to adopt each proposal. Rather, I only argue that they do not solve the problems that arise from strong correctness. Indeed, it would appear that any account where an undecided agent *can* believe $\text{Might } p \wedge \text{Might } \neg p$ will face similar troubles. It is difficult to see how an account would successfully separate *all* instances of indecision from believing $\text{Might } p \wedge \text{Might } \neg p$, without dramatically altering the test conception of “might”. This suggests that the observations about [Willer \(2013\)](#) and [Yalcin \(2018\)](#) generalize, and the issues caused by strong correctness will remain recalcitrant for any plausible view in this tradition. Thus, it does not appear that our account of belief, or the semantics itself, is the locus of the problem. Rather, strong correctness appears to be the culprit.

4 Multiple Content Norms

Evidential, weak, and strong correctness each suffer from different issues, but I suspect that a common cause lies at the root of their respective problems. For truth-conditional frameworks, Wedgwood offered a single content norm in the form of the truth norm. Things like indecision do not come bundled with beliefs and are thus governed by formative norms, making a single content norm perfectly appropriate. The previously considered proposals each attempt the same strategy, modified to fit the test semantics. This strategy, however, seems destined to fail, as the test semantics admits of different kinds of contents that aim to do very different things. Propositional contents aim to represent the way the world is. Non-propositional contents, whatever they do, don’t do that. It is thus little wonder that a single norm is unable to govern such disparate aims. One might also suspect that some of the previous accounts were not beyond repair and could be buttressed with further principles in order to make the right predictions. What follows is an attempt to do right by both of these intuitions.

I propose that adoption of the test semantics should be coupled with the recognition of two content norms. According to these norms, the correctness conditions for every kind of content will not always be the same. Not only will this allow us to distinguish the ways in which different kinds of content can be differently correct or incorrect, but it will also allow us to integrate features of

aforementioned views that were attractive. In particular, the proposal appeals to a condition inspired by weak correctness, as well as a second principle which closely mirrors strong correctness. The difference is that neither is implemented in a vacuum, as each are smaller parts of a larger account. Thus, proponents of the previously considered views may wish to view my proposal as a development of these accounts rather than a direct competitor.

4.1 Locative Correctness and Informational Correctness

Our goal is to draw a distinction between two kinds of content and characterize norms which should govern each. An initially tempting strategy might be to say that propositional contents are governed by something like the truth norm, and non-propositional contents are governed by another. Unfortunately, this won't work. Some non-propositional contents like *Must p* and $p \wedge \text{Might } q$, have propositional entailments, and are, in some sense, beholden to norms sensitive to truth, despite the fact they they are neither true nor false. This suggests that at least some non-propositional contents need to be governed by the same norms as propositional ones.

Given that some non-propositional contents have propositional entailments, we might instead draw our distinction by individuating contents by their propositional entailments. We could do this by saying that some content is correct when its propositional entailments are true. This won't work either, since all contents, propositional and otherwise, entail tautologies, which are propositional. Thus, things like *Might p* will always be correct, since all of their propositional entailments are tautologies. While this strategy doesn't work, it suggests that tautological entailments are the problem. We can avoid these by appeal to *locativity*.

Definition 11: Locativity

φ is locative \Leftrightarrow there exists some contingent proposition $\llbracket \psi \rrbracket$ s.t. $\varphi \Vdash_D \psi$

Some content is locative if and only if that content dynamically entails some propositional content ψ , where $\llbracket \psi \rrbracket$ is contingent.⁴⁰ Locative content individuates certain possibilities from others. Paradigmatically, contingent propositions do just this. They are true at some possible worlds, and false at others. To believe a contingent proposition is to believe that the actual world is a particular way and is not some other way. Similarly, to believe some contingent propositional content is to locate oneself non-trivially in the logical space.

Certain instances of locativity are obvious. Atomic sentences are locative since for any atomic sentence p , $p \Vdash_D p$. The same holds for contingent propositional contents. Alternatively, “might” sentences where the prejacent expresses a contingent proposition are not locative. Sentences like *Might p* do not entail any contingent proposition and thus, do not individuate possibilities from others. Other cases may be less obvious. All contradictions are locative since they entail every formula, contingent propositions included. Tautologies are not

⁴⁰A proposition $\llbracket \varphi \rrbracket$ is contingent $\Leftrightarrow \llbracket \varphi \rrbracket \neq W$ and $\llbracket \varphi \rrbracket \neq \emptyset$.

locative since they are accepted by every state and do not entail any contingent proposition. This makes sense when we consider that tautologies do not individuate any world from any other, and thereby, do not locate in any non-trivial sense. Lastly, “Must” sentences where the prejacent expresses a contingent proposition are locative. $\text{Must } p$, for instance, is locative, since $\text{Must } p \Vdash_D p$. The same holds for “mixed” conjunctions like $p \wedge \text{Might } q$, since $p \wedge \text{Might } q \Vdash_D p$.

Locative content does not coincide with content that is propositional. Many non-propositional contents, most notably $\text{Must } p$, are locative. Similarly, some propositional contents are not locative, e.g., tautologies. Locativity is intended to distinguish those contents which come bundled with commitments about how the world is from those that do not. Once we distinguish locative content from non-locative content, we can characterize how locative content can be correct. Some locative content will be correct when its propositional consequences accurately represent the world. More rigorously:

Locative Correctness: Some locative content φ is locatively correct if and only if for every contingent proposition $[\psi]$, if $\varphi \Vdash_D \psi$ then $@ \in [\psi]$.

Locative content is locatively correct just in case every contingent proposition entailed by that content is true at the actual world. Somewhat loosely, content is correct only when it supports the facts. Locative contents will be locatively correct when they locate the believing agent in the same portion of logical space that contains the actual world. We can project locative correctness to the agential level as well.

Agentially Locative Correctness: An agent, a is locatively correct in believing some locative content φ if and only if φ is locatively correct.⁴¹

Agents will be locatively correct in believing locative contents just in case the contingent propositional commitments incurred by belief in those contents are true. This ensures that correctness in believing at the agential level is determined purely by the content of the belief. It also ensures that locative correctness remains a content norm, and does not encroach into formative territory. Since locative correctness only concerns content, it does not mean that an agent will be correct or incorrect to believe that content; the evidence may or may not support it. However, if that content is locatively correct, an agent will be locatively correct in believing it.

Locative correctness makes the expected and desired prediction about locative contents. Propositional contents that are true (excluding tautologies) are locatively correct. When they are false, they are incorrect. A similar story goes for locative beliefs like $\text{Must } p$, believing which will be locatively correct when p is true at the actual world. If p is false, it is incorrect. All contradictions are locatively incorrect, as they should be. Agents who hold any of these locative beliefs inherit their normative status and are similarly locatively correct/incorrect in believing them.

⁴¹I define correctness at the agential level and the content level. One may also, or instead, choose to define it at the belief level. The difference is trivial since correctness is ultimately decided by the content.

Locative correctness can be seen as something of an extension of the truth norm of belief, but, it accommodates non-truth apt content that comes bundled with truth-apt commitments. Thus, locative belief aims at truth, albeit sometimes indirectly. Rather than requiring the content of the belief to be true, it requires that its contingent propositional consequences be true. Moreover, it captures the ideal of representational belief: to accurately represent the world. It does not, however, tell us anything about contents or beliefs that are not locative. Locative correctness applies only to locative contents and, derivatively, locative beliefs, so all non-locative contents are neither locatively correct, nor locatively incorrect. This makes sense, since they do not attempt to individuate certain possibilities from others. Instead, I propose that non-locative contents are held to a slightly different standard. Call this *informational correctness*.

Informational Correctness: Some content φ is informationally correct if and only if $\{\text{@}\} \models \varphi$.

Some content is informationally correct just in case an idealized state that contains only the actual world supports that content. More intuitively, content is informationally correct if it fits with the facts. Here, we see strong correctness in a maximally idealized form.⁴² The important difference is that informational correctness does not have to work alone, and, as we shall see, interacts with locative correctness in a way that avoids the problems of strong correctness.

Informational correctness provides correctness criteria for all contents, including non-locative contents. For example, if p is true at the actual world, then *Might* p is informationally correct. Given the same facts, *Might* $\neg p$ is informationally incorrect. Tautologies are vindicated since all tautologies are informationally correct. Informational correctness applies to locative contents as well. For instance, non-tautological propositional contents that express true propositions, are informationally correct. False ones are informationally incorrect. Epistemically modalized contents like *Must* p , where p is true, are informationally correct, and when the p is false, incorrect. Contradictions are universally informationally incorrect. Importantly, contents that are locatively incorrect are guaranteed to be informationally incorrect, but not *vice versa*, meaning that locative incorrectness is a special case of informational incorrectness.⁴³ More on this momentarily.

Notice that informational correctness also indirectly concerns truth, but in a different way than locative correctness. Locative correctness demands that any truth-apt entailments of a given locative content be true. Informational correctness instead demands that contents with or without truth-apt entailments ‘fit with’ the facts. Alternatively, we can view informational correctness as ‘alignment with’ the truth. In this sense, informational correctness guides inquiry towards the truth, without representing it. Thus, the truth norm remains in spirit, while not in the letter. We also see echoes of weak and strong

⁴²This formulation is equivalent to the one mentioned in footnote 30.

⁴³Interestingly, locative correctness is not a special case of informational correctness. Suppose that $\text{@} \in \llbracket p \rrbracket$ and $\text{@} \notin \llbracket q \rrbracket$. $p \wedge \text{Might } q$ is locatively correct, but informationally incorrect.

correctness. Locative correctness is inspired by weak correctness, but only concerns contingent consequences. Informational correctness is an idealized form of strong correctness. Unlike before, the two principles interact in a way that allows for finer grained distinctions than were previously available.

Like locative correctness, informational correctness also projects to the agential level.

Agential Informational Correctness: An agent, a is informationally correct in believing some content φ if and only if φ is informationally correct.

To believe some locatively incorrect content is to rule out possibilities that, in some sense, one should not. Thus, we associate locative incorrectness with false contents and false beliefs. Any agent who holds a locatively incorrect belief will not have an information state that is truthful. On the other hand, some content is informationally incorrect when it leaves open possibilities that it, in some sense, should not. However, one can hold a belief that is informationally incorrect while still having an information state that is truthful. Thus, we associate informational incorrectness with contents that fail to align with the truth. Ignorance and indecision, within dynamic frameworks, yield precisely these kinds of beliefs. This is to say that to be locatively incorrect in believing is believing falsely, while being informationally incorrect is failing to believe truly. Believing falsely is a special case of failing to believe truly, so agents that are locatively incorrect are also informationally incorrect.

I say above that these notions of correctness guide what we should *in some sense* believe. Locative correctness captures the sense in which we ought to believe the truth. Informational correctness captures the sense in which our beliefs, propositional and otherwise, should be aligned with the truth. As with any norms concerning content, these are often in conflict with other normative forces that guide our beliefs. Such forces include, but are not limited to, formative norms of belief concerning evidence, inferential mechanisms, etc. My proposal (as well as strong and weak correctness, and any other content norm for that matter) makes no commitments as to the details of these other normative forces, or how to weigh conflicts between them. The upshot is that one can be locatively or informationally correct in holding some belief, yet nonetheless be wrong to believe it, and vice versa. This conflict is equally present in other content governing norms like Wedgewood's thesis.

I take the distinction between locative and informational correctness to be both natural and plausible, and it isn't exactly novel. Traditionally, however, the normative distinction between ignorance and misrepresentation crossed the boundary between content norms and formative norms. Since the test semantics and Heim's conception of belief require that both false belief, as well as indecision involve believing some content, we should expect two content norms in order to distinguish the normative status of each belief. Indeed, semantics that deviate from canonical truth-conditional frameworks are often motivated by the intuition that we ought not to treat contents that express uncertainty in the same way as we treat truth-apt contents. Those that find this attractive should have little trouble extending this reasoning to the norms that govern

content as well.

4.2 Problems Solved

Recognition of locative and informational correctness avoids the second objection to evidential correctness, since neither norm concerns evidence nor strays into formative territory. It also allows us to immediately avoid the homogeneity of error. With the distinction between informational and locative correctness in hand, we are no longer committed to the claim that every incorrect belief is incorrect in the same way. Merely informationally incorrect contents are incorrect in that they fail to be aligned with the truth. It follows that agents who believe such contents will hold beliefs that are not in alignment with, or not “aiming at” the truth. Agents who hold locatively incorrect beliefs similarly fail to appropriately align with the truth, but, in addition, are guaranteed to hold beliefs which misrepresent the world. Such agents are thereby locatively incorrect. This allows us to distinguish the manner in which certain contents, and the agents who believe them, are incorrect. In the absence of the homogeneity of error, the ubiquity of error is appropriately defanged. Numerous contents, and the agents who believe them, are merely informationally incorrect, but this is exactly what we should expect when these beliefs merely fail to align with the truth, without impeding the believing agent’s ability to discover the truth. This is, again, made easy to accept upon appreciation that many of these agents are perfectly correct to believe the contents that they are informationally incorrect in believing.

The distinction between informational and locative incorrectness also fits nicely with the frequent observation that certain undecided agents who hold “might” beliefs with false prejacent are faultless to hold these beliefs. It seems that one is at fault for failing to properly respond to evidence, but this fault is determined by some formative norm of belief. It is not clear, however, that being merely informationally incorrect comes with any fault. Indeed, it would appear that it does not, and that undecided agents can faultlessly hold merely informationally incorrect beliefs. It strikes as plausible to further suggest that one can be considered at fault to hold a false belief. This suggests that agentially locative incorrectness can be associated with some notion of fault, namely the fault one incurs from believing falsely. However, merely informationally incorrect agents do not seem to shoulder the burden of fault. This allows us to maintain the intuition that agents who hold certain “might” beliefs are faultless, while still committing to the claim that the contents of their beliefs are incorrect.⁴⁴ We may thus choose to associate fault with locative incorrectness, but reject that mere informational correctness is associated with fault. This also helps ameliorate some of the sting that may be felt by treating undecided agents

⁴⁴Certain “might” beliefs that are the product of indecision may still be faulty, e.g. suppose that Watson has overwhelming evidence in favor of $\neg m$, but remains recalcitrantly undecided. While Watson may be at fault, this fault is a product of his failure to respond to evidence. Thus, the norm that determines this fault is formative, informational correctness plays no role in determining this fault.

as holding incorrect beliefs. While some of their beliefs will be informationally incorrect, we do not fault them for holding them.⁴⁵

We turn again to Holmes and Watson. Holmes believes $\neg m$ and, therefore, believes that m is false at the actual world. He thus judges Watson’s belief that *Might* m to be informationally incorrect. When Holmes says “No, that’s wrong,” we consider this a judgment of informational incorrectness. More generally, we associate the incorrectness judgments made by disagreeing agents with judgments of informational incorrectness. This is applicable to modal disagreements, as well as straightforward instances of disagreement where some agent holds a propositional belief, and another believes its negation. We can see this by again enriching the scenario to include Lestrade, who believes m . Holmes disagrees with both parties, and is inclined to think that both Watson and Lestrade are informationally incorrect. However, Holmes takes Watson to be merely informationally incorrect, while he believes Lestrade to be locatively, and informationally incorrect. Consequently, Holmes need not believe that Watson’s information state fails to be truthful, but he must believe that Lestrade’s does. This distinction vindicates the intuition that Holmes thinks that Lestrade wrong in a way that Watson is not. This is to say that Holmes recognizes a normative difference between Lestrade’s and Watson’s doxastic positions, despite the fact that they are both informationally incorrect. Should one take on board my earlier suggestion about fault, we may also say that Holmes takes Lestrade to be at fault in virtue of being locatively incorrect, while Watson remains faultless in believing *Might* m , since he is merely informationally incorrect. Importantly, this distinction can be made without appeal to evidence, and only concerns the contents of beliefs. In addition, this is a distinction that non-expert interlocutors are sensitive to. This is supported by data gathered in Khoo (2015), which observes that in instances of modal disagreements, subjects regularly recognize that agents like Holmes need not believe the content of agents like Watson’s belief to be false in order to judge them to be incorrect, (Khoo, 2015, 520).⁴⁶ This suggests that participants can distinguish beliefs that misrepresent the world (and are false) from those that do not represent the world at all but can still be rejected. The difference between locative correctness and informational correctness explains this distinction in addition to vindicating the correctness judgments that motivate the disagreement challenge. The example generalizes, and the same explanations are available across contexts, including eavesdropper cases.

We can tell a similar story with diachronic intrapersonal judgments of incorrectness concerning one’s past beliefs. We can say things like “I used to believe that p might be the case, but I was wrong.” In such cases, we judge ourselves to

⁴⁵These suggestions about fault are merely intended to demonstrate how this proposal fits with intuitions about fault often articulated in the literature. The proposal can be adopted in their absence, should one have differing opinions about how fault is allocated.

⁴⁶Notably, the vignettes in Khoo (2015) are designed such that the evidence available to the characters in the vignettes appear to support the “might” beliefs of the agents that inhabit them. Nonetheless, subjects were still inclined to judge them to be incorrect. This further supports the claim that the normative judgments operative in modal disagreements do not concern formative correctness, but correctness of content.

have been merely informationally incorrect. We can do this while maintaining that, given our evidence at the time, we were correct to believe that p might be the case. We may similarly judge this incorrectness to be faultless. This also allows us to explain the difference between cases like the above, and cases of substantial belief revision, e.g., “I used to believe that p but I was wrong.” Since p is locative, we judge ourselves to have been locatively incorrect, and express that we have since revised our beliefs to reject this incorrect content. We may also judge our past self to have been in some sense at fault for believing falsely.

Individuating informational and locative correctness enables us to solve a number of problems that previous proposals faced. Having two content norms has allowed us to vindicate Holmes’ judgment of Watson’s belief, thus explaining the disagreement. Importantly, it does this in a way that generalizes to all instances of disagreement, modal and otherwise, while simultaneously explaining the manner in which modal disagreements are like, and unlike more straightforward instances of disagreement. However, I have yet to explain how forced decision is avoided. According to my proposal, undecided agents will hold a belief that is informationally incorrect. Since consciously undecided agents should recognize their belief to be informationally incorrect, we must explain what prevents undecided agents like Watson from guessing based solely on the incorrectness of content?

The answer arises from the fact that locative incorrectness is a special case of informational incorrectness. It follows that agents who believe falsely are locatively and informationally incorrect. Undecided agents will hold a belief that is merely informationally incorrect. This is to say that agents who believe falsely aren’t incorrect in an entirely different way than agents who are undecided. Both are informationally incorrect insofar as they hold beliefs that are informationally incorrect, e.g. they are not in alignment with the truth. However, agents who believe falsely hold beliefs that misrepresent the way the world is. Accordingly, agents who are locatively incorrect are incorrect in two ways, whereas agents who are merely informationally incorrect only suffer from one form of incorrectness.⁴⁷ This suggests that locatively incorrect beliefs are, in some sense, worse than merely informationally incorrect beliefs. This captures the intuitive manner in which agents who believe falsely are in an inferior doxastic position than agents who merely fail to believe truly. This also captures the sense in which agents who believe falsely are at fault in a way that agents who merely fail to believe truly are not. This distinction is not only intuitive, but it is antecedently suggested by the test semantics. Agents like Lestrade who hold locatively incorrect beliefs must revise their information state *before* their state can be truthful. Without this, update with the facts will always yield a crash since for any s , if $s \models m$ then $s[\neg m] = \emptyset$. Watson, however, does not need to revise his information before update with $\neg m$ and the update itself will disabuse him of his informationally incorrect belief that Might m . Thus, the test semantics, as well as the norms I’ve proposed, suggest a principle of *doxastic conservativity*.

⁴⁷One might instead say that locatively incorrect agents are wrong in the same way as merely informationally incorrect ones, only worse. Either conception works.

Doxastic Conservativity: It is better to be merely informationally correct than to be locatively incorrect.

According to doxastic conservativity, it is better to hold a belief that is merely informationally incorrect than to hold a belief that is locatively incorrect. This principle explains the aforementioned intuition that Lestrade’s epistemic position is worse than Watson’s. It also explains why undecided agents do not, and should not, guess when they realize that the content of their indecisive belief is informationally incorrect. Suppose a is consciously undecided about p and thus $I_a \models \text{Might } p \wedge \text{Might } \neg p$. a will recognize that the content of their belief is informationally incorrect, however, this will not compel them to guess. In first place, the homogeneity of error no longer holds, and not all incorrect contents are as wrong as false contents. Accordingly, not all incorrectness judgments motivate immediate belief revision. Secondly, two kinds of correctness ensure that there is now substantive risk associated with guessing. With strong correctness, a guess always increased the likelihood of correctness. Once locative and informational incorrectness are distinguished, this no longer holds. If a guesses right then they avoid incorrectness altogether, but if they guess wrong, then they will be locatively and informationally incorrect. Better, then, to remain undecided, and avoid the fault associated with locative incorrectness.⁴⁸ Thus, forced decision is avoided.

Appeal to informational and locative correctness yields the requisite correctness judgements from interlocutors like Holmes. In addition, it avoids all of the problems associated with competing views. It also remains suitably objective, as both informational and locative correctness measure correctness against the facts that populate the actual world. It distinguishes between the incorrect beliefs that are the result of doxastic indecision, and the incorrectness of beliefs that are false. In addition, the separation of these two kinds of correctness is highly plausible and manifests itself in conversations in ways that the laymen is sensitive to. Best of all, extant developments of the test semantics can adopt these correctness conditions for dynamic contents without making other substantive changes to their views. For instance, the lion’s share of [Willer \(2013\)](#), including the entirety of the formal machinery, can remain intact. One need only replace evidential correctness with locative and informational correctness, and one can avoid the problems I’ve discussed. Things are even better for [Lennertz \(2019\)](#), which argues that there are cases of asymmetric disagreement. According to Lennertz, Holmes disagrees with Watson, but Watson does not disagree with Holmes. The arguments made presently are not merely consistent with Lennertz’ position but appear to support it. Note that in virtue of being undecided, Watson does not know whether Holmes is locatively (and thereby informationally) correct or incorrect, since Watson does not know the facts. Meanwhile, Holmes thinks Watson is informationally incorrect and disagrees. If disagreements concerning content come bundled with judgments about correctness, then this seems to vindicate Lennertz’ claim that Watson does not

⁴⁸Importantly, one needn’t take my comments about fault on board in order to make this distinction.

disagree with Holmes, despite the fact that Holmes disagrees with him. It offers further support by bolstering the distinction between asymmetric disagreements like that between Holmes and Watson, and symmetric disagreements like that between Holmes and Lestrade. Holmes thinks Watson is merely informationally incorrect, while Holmes thinks Lestrade is locatively incorrect. Thus, I conclude that proponents of the test semantics should integrate locative and informational correctness into their proposed frameworks.

4.3 Alternative Expressions of Uncertainty

Much of the trouble introduced by strong correctness was the result of the fact that the test semantics, coupled with Heim’s account of belief, ensures that undecided agents will hold certain “might” beliefs. According to my proposal, some of these “might” beliefs are bound to be informationally incorrect. However, an anonymous referee rightly observes that there are alternative ways that one can express uncertainty. According to the proposal so far, such undecided agents should hold the same “might” beliefs. Despite this, Holmes cannot plausibly reject their claims or judge their contents to be incorrect. Observe the discrepancy between the following variations on the scenario between Holmes and Watson.

- (2) a. Watson: Moriarty might be the murderer.
b. Holmes: That’s wrong./You’re wrong./What you said is wrong.
- (3) a. Watson: Moriarty might be the murderer and he might not be.
b. Holmes: That’s wrong./You’re wrong./What you said is wrong./One of those two things is wrong.
- (4) a. Watson: Is Moriarty the murderer?
b. Holmes: # That’s wrong./You’re wrong./What you said is wrong.
- (5) a. Watson: I’m undecided about whether Moriarty is the murderer.
b. Holmes: # That’s wrong./You’re wrong./What you said is wrong.
- (6) a. Watson: I wonder whether Moriarty is the murderer.
b. Holmes: # That’s wrong./You’re wrong./What you said is wrong.

In (2) and (3) Watson expresses his belief that Moriarty might be the murderer. Holmes disagrees because he takes this content to be informationally incorrect, as we have seen. However, in (4)-(6) above, Watson expresses, through different means, that he is undecided as to whether Moriarty is the murderer. Given that Watson is expressing indecision in (4)-(6) and that indecision requires belief in some indecisive content, Holmes should be able to judge this indecisive belief to be informationally incorrect. Despite this fact, Holmes cannot reject Watson’s claims or the beliefs they express in (4)-(6). Moreover, it is less than clear that Holmes disagrees with Watson in the latter three scenarios, and I incur the burden of explaining the discrepancy. This is to say that I must explain why Holmes cannot reject Watson’s utterance, but why he can disagree.

In the problem cases, Holmes cannot reject the content of what Watson says. While each expresses uncertainty in some capacity, (4) is a question, while (5)-(6) are assertions about Watson's own doxastic state. It is unclear, at least from what I've said so far, what it is for the content of a question to be incorrect. Moreover, (5) and (6) are not only assertions, but true ones, insofar as Watson is undecided, and it would not behoove Holmes to reject something true. Nonetheless, according to the account I have provided, there is still a disagreement between Holmes and Watson and Holmes should be licensed to make a normative judgment. I take this to be the correct prediction, and I think informational correctness plays a role in explaining it.

The key here is that rejections and normative judgments where the target of the judgment are not made explicit tend to be parasitic on the content of the sentence previously uttered. Holmes is able to judge the content of Watson's assertions in (2) and (3) above precisely because he judges the content of these assertions, and thereby, the content of Watson's beliefs, to be informationally incorrect. However he cannot do this in (4), because it is a question, nor can he do this with (5) (6) because they are true. This explains why Holmes cannot reject what Watson says; it isn't wrong. However, it remains mysterious why we should think that there is a disagreement.

The answer, I take it, can be gathered from the fact that while Holmes cannot reject what Watson has said, he is perfectly justified in normatively judging the informationally incorrect contents that Watson's speech acts express or otherwise suggest.

- (4') a. Watson: Is Moriarty the murderer?
b. Holmes: You're wrong in thinking he might be.
- (5') a. Watson: I'm undecided about whether Moriarty is the murderer.
b. Holmes: You're wrong in thinking he might be.
- (6') a. Watson: I wonder whether Moriarty is the murderer.
b. Holmes: You're wrong in thinking he might be.

In each case, Holmes is able to deduce that Watson believes that Moriarty might be the murderer, which he judges to be informationally incorrect. Thus, while Holmes cannot negatively normatively judge the content of the initial speech act, he can normatively judge the content of the "might" belief that utterance of that speech act incurs.⁴⁹ Notice that in each case, Holmes' response seems perfectly felicitous and relevant, suggesting that the disagreement was there all along; it simply did not concern the content of Watson's utterances.

⁴⁹Some may object that asking a question does not presuppose that the inquirer is undecided, but instead presuppose that the answer is not established in the common ground. I am sympathetic, but this is no objection to my proposal. If this is indeed the case, then questions do not presuppose ignorance on the part of the hearer, and there is no puzzle to be explained. Still, questions at least seem to suggest that the speaker is undecided. In such cases, Holmes can instead respond to (4a) with "You're wrong *if* you think he might be,"

One need not adopt the account of disagreement I have presupposed in order to appreciate the role that informational and locative correctness play. Even if one takes disagreement to be the action of rejecting a speech act, and thereby rejects the claim that the cases above are genuine disagreements, informational correctness still explains how Holmes is able to felicitously evaluate the contents of Watson's beliefs.⁵⁰ Moreover, it is able to do this in a way that is also able to explain how these evaluations differ from Holmes' evaluation of beliefs he takes to be false, e.g. Lestrade's belief that Moriarty is the murderer.

5 Concluding Remarks

A well-worn metaphor in epistemology analogizes belief and the firing of an arrow. Belief aims at truth in the way that the archer aims at his target. In traditional theories, indecision is the absence of belief, and the arrow is loosed only when a belief is formed. To believe truly is to strike the bullseye. The analogy, however, begins to unravel when applied to previous developments of the test semantics. According to previous accounts, propositional beliefs still aim at truth, but non-propositional beliefs fit poorly into the analogy, as their relationship with truth is indirect. Despite this, they are considered beliefs, proper, and are somehow held to the same standard of correctness as those that directly strive for truth. It is, by my lights, unclear how this fits into an otherwise attractive analogy.

Appeal to locative and informational correctness, in addition to generating the requisite correctness judgments on the part of Holmes, allows the analogy to apply clearly and succinctly to the test semantics. Locative contents indeed aim at the truth, and they hit the target when they are locatively correct. Non-locative beliefs do not fire an arrow, but instead merely take aim at the truth. They are informationally correct when the archer sets his sights on the target. Thus, belief still aims at truth, but not all beliefs loose an arrow.

Metaphors aside, my proposal is fairly simple. The test semantics allows for two distinct kinds of contents. Things like indecision, which are classically treated as absence of belief, now incur belief in content that is not truth-conditional, and so we should not be surprised when a single content norm fails to appropriately govern these two kinds of content. Things are further complexified by the fact that representational contents and non-representational contents have various, often asymmetric, semantic relations with one another. My proposal addresses these issues by insisting that content should be governed by content norms (as opposed to formative ones) and proposing two such norms: informational and locative correctness, and finally characterizing the interactions therebetween. Again, we should not be surprised that problems arise when we try to force formative norms to play a role that they simply cannot play. I take this to be the moral of the story, and I take this moral to be more significant than the details of locative and informational correctness. I

⁵⁰See [Cappelen and Hawthorne \(2009\)](#) for explication of the differences between construal of disagreement as a state versus disagreement as an action.

also expect that this moral can be fruitfully applied to other non-propositional semantic accounts as well. Even if the details differ, we do well to admit of correctness conditions for non-truth-conditional contents that holds them to standards which suit the role they play in our thought and talk.

References

- Aloni, M. (2001) “Quantification Under Conceptual Covers”. (Unpublished doctoral dissertation). University of Amsterdam.
- Beaver, D. (2001) “Presupposition and Assertion in Dynamic Semantics”. CSLI Publications.
- Cappelen, H. Hawthorne, J. (2005) “Relativism and Monadic Truth”. Oxford University Press.
- Dever, J. (2013). “The Revenge of The Semantics-Pragmatics Distinction”. In *Philosophical Perspectives* vol.27 104-144.
- Egan, A. (2007). “Epistemic Modals, Relativism, and Assertion”. In *Philosophical Studies* vol.133, 1-22.
- Egan, A. Weatherson, B. Hawthorne, J. (2005). “Epistemic Modals in Context”. In *Contextualism in Philosophy* Eds. Preyer and Peter. Oxford University Press, 131-169.
- Goldstein, S. (2019). “Generalized Update Semantics”. In *Mind* vol.128(515): 795-835.
- Gillies, A. (2004). “Epistemic Conditionals and Conditional Epistemics”. In *Nous* vol.38(4): 585-616.
- Groenendijk, J. and M. Stokhof (1991a). “Two Theories of Dynamic Semantics”. In *In Logics in AI, Volume 478 of Lectures Notes in Computer Science*. Springer, 55-64.
- Groenendijk, J. and M. Stokhof (1991b). “Dynamic Predicate Logic”. In *Linguistics and Philosophy*. vol.14 (1), 39-100.
- Hawthorne, J. (2007). “Eavesdroppers and Epistemic Modals”. In *Philosophical Issues* vol.17, 92-101.
- Heim, I. (1983). “File Change Semantics and the Familiarity Theory of Definiteness”. In *Meaning, Use, and Interpretation of Language* Eds. R. Bauerle, C. Schwarze, and A. von Stechow. Walter de Gruyter, 164-190.
- Heim, I. (1992). “Presupposition Projection and the Semantics of Attitude verbs”. *Journal of Semantics*, vol.9: 183-221.

- Huvenes, T. (2017). “On Disagreement”. In *Routledge Handbook of Epistemic Contextualism* Ed. Ichikawa, J. Routledge. 272-280.
- James W. (1896). “The Will to Believe”. In *The New World* vol.5, 327-247.
- Kamp, H. (1981). “A Theory of Truth and Semantic Representation”. In *Formal Methods in the Study of Language* Eds. J. A. Groenendijk, T. Janssen, and M. Stokhof. Mathematish Centrum, 277-322.
- Karttunen, L. (1972). “Possible and Must”. In *Syntax and Semantics vol.1* Academic Press. 1-20.
- Khoo, J. (2015). “Modal Disagreements”. *Inquiry* vol.58: 511-534.
- Khoo, J. (2017). “The Disagreement Challenge to Contextualism”. In *Routledge Handbook of Epistemic Contextualism* Ed. Ichikawa, J. Routledge. 257-271.
- Kratzer, A. (1977). “What ‘Must’ and ‘Can’ Must and Can Mean”. In *Words, Worlds, and Contexts: New Approaches to Word Semantics*, ed. Eikmeyer, and Rieser, Berlin: de Gruyter. 38–74.
- Kratzer, A. (1981). “The Notional Category of Modality”. In *Linguistics and Philosophy*, vol. 1, 337-355.
- Lennertz, B. (2019). “Might-Beliefs and Asymmetric Disagreement”. *Synthese* vol.196: 4775–4805.
- Lennertz, B. (2020). “Evaluating the Multiple Proposition Strategy”. *Ratio* vol.33(3): 163-172.
- Lewis, D. (1980). “Index, Context, and Content”. In *Philosophy and Grammar*, Eds. Kanger, S. and Ohman, S. Dordrecht:Reidel, 79-100.
- MacFarlane, J. (2011). “Epistemic Modals are Assessment Sensitive”. In *Epistemic Modality* Eds. Egan and Weatherson. Oxford University Press, 144-178.
- MacFarlane, J. (2014). *Assessment Sensitivity*. Oxford University Press.
- Mandelkern, M. (2020). “Dynamic Non-Classicality”. In *Australian Journal of Philosophy* vol.98 (2) 382-392.
- Nouwen, R. Brasoveanu, A. van Eijck, A. Visser, A. (2016). “Dynamic Semantics”. In *The Stanford Encyclopedia of Philosophy* Ed. Zalta. <<https://plato.stanford.edu/entries/dynamic-semantics/>>. Accessed 1/2/2020.
- Plunkett, D., Sundell, T. (2013). “Disagreement and the Semantics of Normative and Evaluative Terms”. In *Philosopher’s Imprint* vol.13(23): 1-37.
- Rothschild, J. Yalcin, S. (2017). “On the Dynamics of Conversation”. In *Nous* vol.52 (2), 24-48.

- Stalnaker, R. (1978). “Assertion”. In *Syntax and Semantics 9: Pragmatics*, Ed. Peter Cole. 315–332. Academic Press.
- Stalnaker, R. (1998). “On the Representation of Context”. In *Journal of Logic, Language, and Information* vol.7, 3-19.
- van Benthem, J. (1989). “Semantic Parallels in Natural Language and Computation”. In In Ebbinghaus, H. D., Fernandez-Prida, J., Garrido, M., Lascar, D. and Artalejo, M. R. (eds.) *Logic Colloquium '87: Proceedings of the Colloquium held in Granada, Spain*. Volume 129 of Studies in Logic and the Foundations of Mathematics. 337–75. Amsterdam: North-Holland.
- van der Does, J., Groeneveld, W. Veltman, F. (1997). “An Update on “Might””. In *Journal of Logic, Language, and Information* 6(4): 361–380.
- Veltman, F. (1996). “Defaults in Update Semantics”. In *Journal of Philosophical Logic* vol.25, 221-261.
- von Fintel, K. Gillies, A. 2007. “An Opinionated Guide to Epistemic Modals”. In T. Szabó Gendler and J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 2). Oxford: Oxford University Press.
- von Fintel, K. Gillies, A. 2010. “Must... Stay... Strong!”. In *Natural Language Semantics* vol.18 (4) 351-383.
- von Fintel, K. Gillies, A. 2011. “Might made Right”. In *Epistemic Modality* Eds. Egan and Weatherson. Oxford university Press, 108-130.
- Wedgewood, R. (2002). “The Aim of Belief”. In *Philosophical Perspectives* vol.16 267-297.
- Williams, B. (1973). “Deciding to Believe”. In *Problems of the Self*. Cambridge University Press. 136-151.
- Willer, M. (2013). “New Dynamics for Epistemic Modality”. In *Philosophical Review* vol.122 (1) 45-92.
- Willer, M. (2015). “An Update on Epistemic Modals”. In *Journal of Philosophical Logic* vol.44 (6) 835-849.
- Yalcin, S. (2007). “Epistemic Modals”. In *Mind* vol.116 (464), 983-1026.
- Yalcin, S. (2011). “Nonfactualism About Epistemic Modality”. In *Epistemic Modality* Eds. Egan and Weatherson. Oxford university Press, 295-332.
- Yalcin, S. (2012a). “Introductory Notes on Dynamic Semantics”. In *Routledge Companion to the Philosophy of Language* Eds. Fara and Russell. Routledge, 253-279.
- Yalcin, S. (2012b) “Bayesian Expressivism”. In *Proceedings of the Aristotelian Society* vol.112 (2), 123–160.

Yalcin, S. (2015) “Epistemic Modality *de re*”. In *Ergo* vol.2 (19), 475-527.

Yalcin, S. (2018) “Belief as Question Sensitive”. In *Philosophy and Phenomenological Research* XCVII (1), 23-47.