

Assertions of Clarity & Raising Awareness

Abstract

Barker and Taranto (2003) introduce the “paradox of asserting clarity,” arguing that any true assertion of the form *It is clear that p* should be necessarily uninformative. Assuming that the essential effect of assertions is to inform, it is therefore puzzling why any speaker would bother asserting that a proposition is clear. I present a novel account of the semantics of *clear* that predicts assertions of clarity to be straightforwardly informative in some contexts, thereby partially resolving this paradox. However, on this account there remain cases in which clarity assertions are not obviously informative. I argue that in such cases, an assertion of the clarity of *p* may play the role of raising awareness of *p* or of evidence supporting *p*. Raising awareness of *p* may, in turn, have a number of downstream consequences; in particular, raising an addressee’s awareness of an issue may have crucial effects on the addressee’s resolution to decision problems. The effect of clarity assertions on agents’ awareness states is formalized using a model of awareness in discourse inspired by Franke and de Jager (2007), de Jager (2009), and Franke and de Jager (2011).

1 Introduction

Barker and Taranto (2003) argue that any true assertion of the form *It is clear that p* is necessarily uninformative. To illustrate this argument, consider the following scenario. Little Jack Horner is in San Francisco and is looking for a good place to get pastries. He enlists the help of his friend Little Bo Peep. They come across two bakeries: Tartine and Arizmendi. When Jack and Bo approach the bakeries, there is a long line outside Tartine, but there is no line outside Arizmendi. Bo utters the following:

- (1) BO: It’s clear that there is a wait at Tartine.

Intuitively, Bo’s utterance in (1) is true in the context described. I take this situation as a starting point for understanding the truth conditions of clarity statements.

First, for Bo’s utterance to be true, she must possess evidence that there is a wait at Tartine. In the situation described above, her evidence consists in the line of people waiting outside Tartine. Because she lacks evidence that there is a wait at Arizmendi, Bo cannot felicitously assert (2) in the context described above:

- (2) BO: # It’s clear that there is a wait at Arizmendi.

Bo’s utterance in (2) remains infelicitous even if Bo has private evidence that Arizmendi has a wait. Perhaps Bo’s friend Humpty Dumpty is inside Arizmendi and has sent a text message to Bo saying that there is a wait. Still, it would be odd for Bo to assert to Jack that it is *clear* that there is a wait at Arizmendi unless Jack has access to the relevant evidence as well. Based on similar considerations, Barker and Taranto offer the following initial hypothesis about the truth conditions of clarity statements: *It is clear that p* is true if and only if there exists some publicly available evidence, i.e. evidence available to all discourse participants, that supports the conclusion that *p*.

Now suppose that the clarity of *p* is truthfully asserted in some context. It follows that there is some publicly available evidence in this context that supports the conclusion that *p*. Assuming that all discourse participants are rational, they should all conclude *p* on the basis of the publicly available evidence. But if all discourse participants have already concluded *p*, and can conclude that all other discourse participants have concluded *p*, then the assertion that *p* is clear does not

seem to contribute any new information in the discourse. In other words, any true assertion of the clarity of p appears to be uninformative. This result is paradoxical on the standard assumption, going back to Stalnaker (1978), that the purpose of assertion is to inform and that uninformative assertions are therefore infelicitous. Following this line of reasoning, we conclude that the clarity of a proposition p can never be felicitously asserted, since such an assertion is either uninformative or false. Yet (1) shows that clarity *can* be felicitously asserted, so there must be some error in the preceding argument.

A number of authors have previously tackled this puzzle (Barker and Taranto 2003; Taranto 2006; Bronnikov 2008; Barker 2009; Wolf and Cohen 2011; Wolf 2014), each by developing a lexical semantics for *clear* according to which, it is argued, clarity assertions may be both true and informative. In part, I also follow this strategy in that I offer a new proposal for the meaning of *clear* that allows for clarity assertions to be both truthful and straightforwardly informative in certain contexts. This new proposal for the semantics for *clear* takes inspiration from genericity-based accounts of predicates of personal taste.

However, I also argue that developing the correct semantics for *clear* is insufficient to explain the function of clarity statements in all contexts. In contexts in which assertions of clarity are not straightforwardly informative, I explain these assertions' felicity via utterances' ability to manipulate an addressee's attention or awareness. In turn, raising awareness of a proposition may have important consequences for the structure of the discourse and for some relevant decision problem. This approach to understanding the function of clarity assertions does not rely on any unique properties of clarity assertions *per se*, but rather addresses the more general question of what purpose, if any, seemingly uninformative assertions can have. Nonetheless, I take this proposal to be a valid resolution to the paradox of asserting clarity in that resolving this paradox requires the identification of some function that clarity assertions can serve, whether or not this is a function that is *only* served by clarity assertions. However, in the conclusion of this paper, I offer thoughts as to why clarity assertions may be preferred over other utterances for awareness-raising uses in some contexts.

The rest of the paper is organized as follows. In the remainder of this section, I flesh out the paradox of asserting clarity and present some initial desiderata for the semantics of clarity statements. In §2, I discuss previous proposals for resolving the paradox. In §3, I propose a new semantics for clarity statements that takes inspiration from genericity-based analyses of predicates of personal taste. I then show how, on this new proposal, clarity statements may be both true and informative. Nonetheless, there remain cases in which true clarity assertions are not obviously informative. In §4, I address such cases by illustrating how clarity assertions can serve the purpose of manipulating an addressee's state of awareness or attention. Finally, I conclude in §5 with some discussion of why a speaker might choose to raise awareness of an issue with an assertion of clarity, as opposed to alternative expressions.

1.1 The Paradox

Before we consider candidate resolutions to Barker and Taranto's paradox, it is worth trying to restate the paradox somewhat more formally to have a clearer sense of what, precisely, must be resolved. Note that the original formulation of the paradox relies on informal, intuitive notions regarding whether some proposition can be rationally inferred from some body of evidence. The restatement below also relies on these intuitive notions, and as a consequence is not fully formally precise. But note that the purpose of this restatement is not to show that the paradox of asserting clarity is, in fact, a logical paradox. In the end, we will want to show that clarity assertions are not logically paradoxical. Rather, it is to make more apparent why our intuitions might lead us to

believe that clarity statements should be puzzling in the first place.

Building upon work such as Stalnaker (1978), Heim (1982), and Stalnaker (2002), we define the common ground as follows:¹

- (3) Let I be the set of all interlocutors in some discourse. Then the COMMON GROUND CG is the set of propositions p such that for all $i \in I$, i believes p and i believes that for all $i' \in I$, i' believes p , and i believes that for all $i' \in I$, i' believes that for all $i'' \in I$, i'' believes p , etc.

Taking the intersection of these propositions gives us the context set, the set of worlds consistent with the common ground:

- (4) Let CG be a common ground as defined in (3). Then the CONTEXT SET C is $\cap CG$.

An assertion alters the common ground as follows. Suppose that a discourse participant asserts A such that the content of A is p . If A is accepted, p enters the common ground.² Let C and C' denote the context set of a discourse before and after the acceptance of A , respectively. We have $C' = C \cap p$. A is informative if and only if $C' \subset C$.³

Turning to clarity assertions, suppose that in a context with context set C , a speaker felicitously asserts that some proposition p is clear. Let $clear(p)$ denote the content of the speaker's assertion. If clarity assertions are truly paradoxical in the way described by Barker and Taranto, we should find that so long as the speaker's assertion is felicitous, it is guaranteed that $C = C \cap clear(p)$. In other words, if an assertion of the clarity of p is felicitous, then the common ground must have already entailed $clear(p)$ before the assertion was made. Why should we think that clarity assertions are paradoxical in this way? We can motivate Barker and Taranto's paradox by adopting the following assumptions:

- (5) a. An assertion of $clear(p)$ is true if and only if there exists some set of propositions $e \subseteq CG$ (the evidence) such that any agent who believes e can rationally infer that p .
 b. All discourse participants are rational and logically omniscient.
 c. It is common ground that all discourse participants are rational and logically omniscient.

(5a) simply gives the initial truth conditions suggested above for $clear(p)$. These truth conditions can be broken down into two parts: (i) e is in the common ground; and (ii) any agent who believes e can rationally infer that p . To motivate the paradox, we must show that both (i) and (ii) must already be in the common ground before $clear(p)$ is felicitously asserted. Of course, whenever

¹The common ground is defined here in terms of the common beliefs of all discourse participants. However, as it is usually conceived, the common ground may contain propositions that are not common belief. Participants in a discourse may suppose that some propositions are true for the sake of the conversation, without being committed to those propositions being true in the actual world. Such propositions are usually included in the common ground, although I ignore this detail for the sake of simplicity.

²Much recent work, such as Davis (2009); Farkas and Bruce (2010); Ginzburg (2012); Gunlogson (2001, 2008); Lauer (2013) and Malamud and Stephenson (2015), offers a more nuanced perspective on the conventional discourse effects of assertions. In general, on this view the essential discourse effect of assertion A in terms of a speaker commitment to p , the content of A , and a proposal to update the CG with p . Only when A is accepted by discourse participants is p added to the CG . Since I am concerned primarily with cases in which clarity assertions are accepted, I do not engage with models of discourse which represent proposals to update the CG .

³Using the definitions given here, it is possible for an assertion A to be uninformative even if $\llbracket A \rrbracket$ is not in the common ground. For example, suppose both p and q are part of the common ground, but the proposition $p \cap q$ is not. If $\llbracket A \rrbracket = p \cap q$, asserting A would add $p \cap q$ to the common ground, but not have any effect on the context set. Thus, asserting A would be uninformative. To avoid this consequence, we could further assume that the common ground is closed under entailment.

$clear(p)$ is felicitously asserted, (i) and (ii) hold, so we can alternatively view our task as showing that whenever (i) and (ii) are true, (i) and (ii) are common ground.

First, it follows from (i) that e is already in the common ground. From the definition of the common ground in (3), that e is in the common ground is also in the common ground. Thus, (i) must be in the common ground before $clear(p)$ can be felicitously asserted. Next, it follows from (ii) that e is sufficient evidence for any agent to rationally infer p . By (5b), all discourse participants should have inferred p from the evidence e . Assuming a form of introspection according to which rational agents can always recognize that a particular inference is rational, all discourse participants should also conclude that e is sufficient evidence to rationally infer p . Finally, by (5c), all discourse participants can recursively reason about what the other discourse participants have inferred on the basis of e . In particular, all discourse participants should recognize that all other discourse participants have inferred p , that e is sufficient to infer p , etc. This ensures that (ii), the proposition that any agent who believes e can rationally infer that p , is in the common ground. Thus, by the assumptions in (5), we find that whenever (i) and (ii) are true, (i) and (ii) are common ground. Alternatively, $clear(p)$ must already be in the common ground whenever $clear(p)$ is truthfully asserted.

One may object that our initial observations regarding the felicity of clarity assertions do not justify (5a). At best, all we have seen is that the existence of publicly available evidence supporting p is a necessary condition for the clarity of p , but perhaps it is not a sufficient condition for the clarity of p . Suppose we change (5a) to the following:

- (5) a'. If an assertion of $clear(p)$ is true, then there exists some set of propositions $e \subseteq CG$ (the evidence) such that any agent who believes e can rationally infer that p .

On this modification, an assertion of the clarity of p still requires the presence of some evidence in the common ground from which p can be rationally inferred. But the truth of $clear(p)$ may also depend on other factors, which we have not yet specified and which may not be common ground.

However, even with the modified assumption in (5a'), assertions of clarity are puzzling in the following way. If $clear(p)$ has been felicitously asserted, we know by (5a') that all discourse participants believe e and that all discourse participants can rationally infer p . By (5b), it follows that all discourse participants have inferred p . Finally, by (5c), it follows that all discourse participants have inferred that all other discourse participants have inferred p . This ensures that p is in the common ground before the clarity of p is asserted. While this falls short of showing that the assertion of $clear(p)$ is uninformative, is mysterious why anyone would bother to assert $clear(p)$ if p is already in the common ground.

To resolve the paradox, we will have to abandon at least one of the assumptions in (5). But not all ways of rejecting these assumptions will fully resolve the paradox. In particular, if we only modify (5a) by assuming (5a'), we must still answer the question of what the effect of asserting $clear(p)$ is when p is already in the common ground.

1.2 Initial Desiderata for a Semantics of Clarity

The above discussion of the paradox of asserting clarity relied on a first-pass hypothesis regarding the meaning of clarity statements. But in order to resolve this paradox, it is necessary to have a more detailed understanding of the truth conditions of $clear$. Here, I lay out the phenomena that any semantic analysis of clarity statements must account for.

The paradox of asserting clarity is illustrated above with an assertion of the form *It is clear that p* (SIMPLE CLARITY). However, other types of clarity statements are possible, which make

explicit *to whom* a proposition is clear (PERSONAL CLARITY) or *from what evidence* the clarity of a proposition follows (EVIDENTIAL CLARITY):⁴

- (6) a. Simple clarity
It's clear that there is a wait at Tartine.
- b. Personal clarity
It's clear to me that there is a wait at Tartine.
- c. Evidential clarity
It's clear from the line outside that there is a wait at Tartine.
- d. Personal-evidential clarity
It's clear to me from the line outside that there is a wait at Tartine.

It is also useful to establish terminology for the arguments appearing in these expressions. Suppose we have a personal-evidential clarity expression of the form *It is clear to x from e that p*. Call *p* the CONCLUSION or PREJACENT of a clarity statement, *x* the EXPERIENCER, and *e* the EVIDENCE. An ideal semantics for *clear* would provide a meaning for each type of clarity expression and elucidate the relationship between these different expressions.

The paradox of asserting clarity is generally taken only to apply to assertions of simple clarity. Assertions of personal clarity and personal-evidential clarity are not taken to be problematic because they do not appear to be uninformative in the same way as assertions of simple clarity. Recall that the paradox relied on the notion that the evidence supporting a clarity assertion must be possessed by all discourse participants. In contrast, it is not hard to construct examples of personal and personal-evidential assertions of clarity that rely on private evidence. Consider (7):

- (7) BO: It's clear to me that Tartine has the best pastries in town, but I know you've never been there.

The evidence for Bo's assertion is (presumably) her experiences with pastries from Tartine and other bakeries, and this evidence is not shared by her interlocutor. Therefore, there is no expectation that her interlocutor can make any conclusions regarding the quality of Tartine's pastries, and (7) may be straightforwardly informative.

There are important relationships between the conclusion, experiencer, and evidence of clarity expressions that should be captured by our semantics for *clear*. First, the experiencer must have knowledge of the evidence, as shown by the infelicity of (8):

- (8) # It is clear to Bo from the line outside that there is a wait at Tartine, but Bo does not know whether there is a line outside Tartine.

Second, as noted above, the evidence must support the conclusion. Thus, it is infelicitous to assert the clarity of *p* on the basis of evidence that is completely irrelevant to *p*:

- (9) # It is clear from the fact that nine is divisible by three that there is a wait at Tartine.

Moreover, even if the evidence is relevant to *p*, the clarity of *p* cannot be asserted if this evidence is insufficient to conclude *p*. For example, suppose that Bo is in the running for a promotion at work and we have evidence sufficient to conclude that it is clear that Bo will receive the promotion.

⁴For the most part, I ignore adverbial clarity expressions, e.g. *Clearly, there is a wait at Tartine*, but assume that a sentence of the form *Clearly p* is equivalent to *It is clear that p*. An alternative view is offered by Wolf and Cohen (2011), who argue that adverbial clarity statements require a different analysis than non-adverbial clarity statements. Building on Piñon (2006), Wolf and Cohen claim that adverbial clarity statements are speaker-oriented, so *Clearly p* is roughly equivalent to *It is clear to me that p*. See Barker (2011) for some criticisms of Wolf and Cohen's treatment of *clearly*.

Suppose that this evidence includes the fact that Bo was employee of the month last month, but that this fact alone is not sufficient to conclude that she will be promoted. In this case, (10) is infelicitous.

- (10) # It is clear from the fact that Bo was employee of the month last month that she will be promoted.

Thus, any semantics for *clear* must capture the fact that the clarity of p from e requires that e is sufficient to conclude p . I refer to this as the EVIDENTIAL SUFFICIENCY REQUIREMENT.

If an experiencer must be in possession of the evidence, and this evidence must be sufficient to conclude p , then we should also expect the experiencer to believe the conclusion of a clarity expression. At least as a first pass, this seems correct:

- (11) # It is clear to Bo that there is a wait at Tartine, but Bo does not believe there is a wait at Tartine.

As we will see below, on some (but not all) analyses of clarity expressions, the clarity of p to x entails that x believes p . Still, all analyses must find some way of explaining the infelicity of examples like (11).

Barker and Taranto (2003) and Barker (2009) note that *clear* is vague, which can be confirmed by showing that *clear* exhibits the properties standardly associated with vague predicates: (i) the truth conditions of sentences containing vague predicates are contextually variable; (ii) vague predicates have “borderline cases” in which it is difficult to say whether the predicate truly applies or not; and (iii) vague predicates give rise to Sorites paradoxes (Kennedy 2007).

The contextual variability of vague predicates is usually tied to the fact that the interpretation of a vague predicate depends upon a contextually relevant comparison class. To see such contextual variability in the case of *clear*, consider the following:

- (12) It is clear that P does not equal NP.

It is currently unknown whether P equals NP, although most computer scientists believe that P does not equal NP (Rosenberg 2012). In a discussion about proven mathematical results, (12) would likely be judged false on the basis of the fact that it is unknown whether P equals NP. But in a discourse in which the P versus NP problem is compared to other unanswered mathematical questions, we might judge (12) to be true, since many experts share the strong intuition that it is true.

Borderline cases and the Sorites paradox can be illustrated with the following example. Suppose we want to determine whether a coin is fair or not. After flipping the coin once, it comes up heads. In this context, (13) is false.

- (13) It is clear that the coin is unfair.

We then flip the coin ninety-nine more times, and it comes up heads each time. After these coin flips, (13) is judged to be true. The more interesting question is whether (13) is true at some intermediate point, e.g. is it clear that the coin is unfair after it has been flipped 10 times? At such a borderline case, it is difficult to say whether the proposition that the coin is unfair is clear or not. A Sorites paradox for *clear* can be constructed by first taking as a premise that if it is clear that a coin is unfair after coming up heads x times in a row, then it is clear that coin is unfair after coming up heads $x - 1$ times in a row.⁵ Next, we take as a premise that it is clear that a coin is unfair after coming up heads 100 times in a row. We can then conclude that it is clear that a coin

⁵See Barker (2009) for an alternative Sorites paradox for *clear*.

is unfair after coming up heads just once. While both premises are plausible, the conclusion is not. Thus, *clear* has the standard properties associated with vague predicates.⁶

Having set out some initial desiderata for the semantics of *clear*, in the next section I consider several previous proposals for the meaning of clarity statements. In each case, my primary concern is with and how well these proposals are able to resolve the paradox of asserting clarity, but I also comment on how well they address the properties of *clear* noted in this section.

2 Previous Proposals

2.1 Barker and Taranto (2003)

Barker and Taranto (2003) propose that a proposition p is clear to an agent so long as the agent's degree of belief or credence in p exceeds some contextually relevant threshold.⁷ Simple clarity is derived from personal clarity by taking the implicit experiencer argument to default to the set of all discourse participants. Let $\mathbf{d}(c)(\llbracket clear \rrbracket)$ be a function from contexts c to degrees of likelihood. Then, we have the following definitions for personal and simple clarity:⁸

- (15) a. *It is clear to x that p* is true in a context c iff the maximal degree to which x judges that p is likely to be true is at least as great as $\mathbf{d}(c)(\llbracket clear \rrbracket)$.
- b. *It is clear that p* is true in a context c iff the maximal degree to which all discourse participants judge that p is likely to be true is at least as great as $\mathbf{d}(c)(\llbracket clear \rrbracket)$.

Barker and Taranto attempt to resolve the paradox of asserting clarity by leveraging the distinction between mutual belief in a proposition p , which only requires that all discourse participants believe p , and common belief as defined above. Given the semantics in (15), a proposition p is clear in a context so long as it is mutually believed by all discourse participants. Since mutual belief is a weaker notion than common belief, the truth of *clear*(p) does not ensure that either p or *clear*(p) is in the common ground. However, once the clarity of p is asserted, all discourse participants know that p is mutually believed, and they know that all other discourse participants know that p is mutually believed, etc. Thus, an assertion of *clear*(p) converts the mutual belief in p into common belief in p .

Now consider this proposal in terms of the assumptions laid out in (5). Although their truth conditions do not make reference to the evidence supporting clarity statements, Barker and Taranto's prose indicates that they accept (5a'), which states that the evidence is in the common ground and that any rational agent in possession of this evidence can conclude the prejacent. They

⁶ Interestingly, *clear* patterns with maximum-standard absolute adjective (Kennedy and McNally 2005).

- (14) a. It is 100% clear/??unclear that Mindy is a doctor.
 b. It is completely clear/??unclear that Mindy is a doctor.
 c. It is perfectly clear/??unclear that Mindy is a doctor.
 d. It is fully clear/??unclear that Mindy is a doctor.

According to Kennedy and McNally (2005), for some maximum-standard absolute adjective a to be truly predicated of an entity x , x must be maximally a , i.e. exhibit a maximum degree of the scale denoted by a . From this, Kennedy and McNally claim that maximum-standard absolute adjectives should not exhibit vagueness. Exceptions to this generalization, such as *bald*, have been noted (Kennedy, 2007, fn. 30), and it seems that *clear* is another exception.

⁷The proposals put forward in Taranto (2005) and Taranto (2006) are largely similar to that offered in Barker and Taranto (2003). For this reason, I do not discuss them separately.

⁸Absent from Barker and Taranto's proposal is any discussion of the role of evidence in the truth conditions for clarity, the evidential sufficiency requirement, or a proposal for the truth conditions of evidential clarity statements.

also state their commitment to (5b), which states that all discourse participants are rational and logically omniscient. On the basis of these two assumptions, it follows that whenever $clear(p)$ is truthfully asserted, p must be a mutual belief of all discourse participants.

For their proposed solution to work, Barker and Taranto must reject (5c), which states that it is common ground that all discourse participants are rational and logically omniscient. By rejecting (5c), it is possible for p to be a mutual belief but not a common belief. Note, however, that Barker and Taranto offer no justification for their (implicit) rejection of (5c), but their commitment to (5b). A plausible motivation for rejecting (5c) is that it is psychologically unrealistic to assume perfect rationality and logical omniscience on the part of all discourse participants. But if we use this reasoning to reject (5c), we can use the same reasoning to reject (5b). Indeed, this type of reasoning will be used to motivate a rejection of both (5b) and (5c) below. But Barker and Taranto's stance, in which (5b) is accepted, but (5c) is rejected, is not well-motivated.

2.2 Barker (2009)

Barker (2009) offers a revised analysis of clarity that moves away from the belief-based account of Barker and Taranto (2003). This move is motivated by several arguments that belief is neither necessary nor sufficient for clarity. First, Barker notes that even in a discourse in which all participants strongly believe that God exists, they may still judge (16) to be false.

(16) It is clear that God exists.

Thus, mutual belief in p is not a sufficient condition for the clarity of p . Next, Barker claims that a speaker could assert (17) while holding a private belief that there is life on Mars.

(17) It is clear that Mars is barren of life.

According to Barker, a speaker who asserts (17) claims that it is rational, given the evidence, to conclude that Mars cannot support life. It would be very odd for a speaker to assert this without actually believing that Mars is barren of life, but Barker maintains that a speaker who asserts (17) may hold an irrational belief that there is life on Mars. Thus, according to Barker, belief in p is not a necessary condition for the clarity of p .

Barker offers a new set of truth conditions for clarity claims, based on the idea that a proposition is clear so long as it is well-justified by the common ground. The notion of how well a proposition is justified by the common ground is formalized in the style of Kratzerian analyses of modality (Kratzer 1981, 1991). Let C be the context set, and let $g : \mathcal{W} \rightarrow \wp(\wp(\mathcal{W}))$ be a stereotypical ordering source from worlds w to assumptions about the normal course of events in w . At a world w , the ordering source $g(w)$ imposes a partial order on worlds $\leq_{g(w)}$ such that $w' \leq_{g(w)} w''$ if and only if $\{p \in g(w) | w'' \in p\} \subseteq \{p \in g(w) | w' \in p\}$. Next, let $\mu(w)$ be a function mapping worlds to degrees such that $\mu(w') \leq \mu(w'')$ only if $w' \leq_{g(w)} w''$. Finally, let $\mathbf{d}(w)(\llbracket clear \rrbracket)$ be a function that maps each world w onto the necessary degree of justification required for a proposition to be clear in w . Simple clarity is defined as follows:⁹

(18) *It is clear that p* is true at a world w iff $\forall w' \in C (\mu(w') \leq \mathbf{d}(w)(\llbracket clear \rrbracket) \rightarrow w' \in p)$.

⁹Barker does not provide a formal definition of personal clarity, but writes that personal clarity may rely on evidence that is possessed by the experiencer, but is not part of the common ground, and that personal clarity involves a personal standard for clarity, rather than one that is in the common ground. Thus, for a personal clarity statement to be true, all worlds consistent with experiencer's body of evidence must be at least as normal as the experiencer's personal standard for clarity. Barker, though acknowledging the existence of evidential clarity statements, does not comment on how the truth conditions for evidential clarity expressions differ from (18).

That is, a proposition is clear so long as it is true in all of the most normal worlds consistent with the common ground. Barker takes both the evidence and the contents of $g(w)$ to be part of the common ground. Thus, the common ground settles every proposition's degree of justification up to a positive monotonic transformation. However, this does not settle whether a proposition is clear, as Barker contends that discourse participants may be in doubt about the contextual standard $\mathbf{d}(w)(\llbracket clear \rrbracket)$. Assertions of clarity provide information about this standard, since following a felicitous assertion of the clarity of p , $\mathbf{d}(w)(\llbracket clear \rrbracket)$ can be no greater than the degree of normality of most normal non- p world.

In addition, Barker takes the value of $\mathbf{d}(w)(\llbracket clear \rrbracket)$ to play a key role in which propositions may enter the common ground. We can say that a proposition p has a degree of justification \mathbf{d}_p when \mathbf{d}_p is the smallest value such for all $w \in p$, $\mu(w) \leq \mathbf{d}_p$. Then, only propositions whose degrees of justification are equal to or exceed $\mathbf{d}(w)(\llbracket clear \rrbracket)$ may be part of the common ground. A speaker who asserts the clarity of p ensures that p and all other propositions as well-justified as p may enter in the common ground. In this way, the proposal of Barker (2009) maintains Barker and Taranto's (2003) idea that asserting the clarity of p is a way to ensure that p enters the common ground. However, in the revised proposal, this is accomplished indirectly by setting the standard $\mathbf{d}(w)(\llbracket clear \rrbracket)$.¹⁰

To provide an illustration of how Barker's proposal would work, suppose Bo and Jack are in a context such that for all $w \in C$, $g(w)$ contains only two propositions: p , the proposition that if there is a line outside a bakery, there is a wait in the bakery, and q , the proposition that the croissants at Tartine are fresher than those at Arizmendi. For every world $w \in p \cap q$, the value of $\mu(w)$ will be the same. Let \mathbf{d}_{pq} denote this value. Likewise, the value of $\mu(w)$ will be the same for every world $w \in p \cap (\mathcal{W} \setminus q)$. Let \mathbf{d}_p denote this value. We have $\mathbf{d}_{pq} < \mathbf{d}_p$. Now suppose that Bo and Jack see the line outside Tartine and that the proposition that there is a line outside Tartine enters the common ground. Bo then utters that it is clear there is a wait at Tartine. The immediate effect of this utterance is to set $\mathbf{d}(w)(\llbracket clear \rrbracket)$ to a value that is no greater than \mathbf{d}_p . As a result, p , and everything entailed by p , enters the common ground.¹¹

Barker's approach can be viewed as a rejection of (5a), the assumption that if an assertion of $clear(p)$ is felicitous, there exists evidence in the common ground sufficient for any agent to conclude p . Barker makes the follow, more nuanced assumption:

- (5) a''. If an assertion of $clear(p)$ is felicitous, there must be some set of propositions $e \subseteq CG$ (the evidence) and some degree \mathbf{d} such that any agent who believes e can rationally infer that p is justified to at least degree \mathbf{d} .

If we assume (5a''), along with (5b) and (5c), it follows that in any context in which it is felicitous to assert $clear(p)$, it is common ground that p is justified to degree \mathbf{d} . But this does not ensure

¹⁰Although Barker (2009) argues the only function of clarity assertions is to set the standard $\mathbf{d}(w)(\llbracket clear \rrbracket)$, Barker (2011) proposes another function for these assertions. According to Barker (2011), interlocutors may disagree about the ordering source $g(w)$, a position that is intended to account for cases discussed by Wolf and Cohen (2011) in which interlocutors conclude that contradictory propositions are clear on the basis of the same bodies of evidence. If interlocutors shared the ordering source $g(w)$, as assumed by Barker (2009), such situations would be impossible. On the revised view presented in Barker (2011), clarity assertions may serve to inform discourse participants about the appropriate ordering source for the normality of worlds in addition to resolving the contextual standard for clarity $\mathbf{d}(w)(\llbracket clear \rrbracket)$.

¹¹Of course, the value of $\mu(w)$ will also be the same for every world $w \in q \cap (\mathcal{W} \setminus p)$. We can denote this value with \mathbf{d}_q . Worlds in $p \cap \mathcal{W} \setminus q$ and $q \cap \mathcal{W} \setminus p$ are incomparable with respect to $\leq_{g(w)}$, so depending on the choice of μ , we may either have $\mathbf{d}_p > \mathbf{d}_q$ or $\mathbf{d}_p < \mathbf{d}_q$. Thus, under some choices of μ , the assertion of the clarity of p will also inadvertently cause q and everything entailed by q to enter the common ground. To avoid this consequence, we can require that a proposition enters the common ground so long as it is true at every world w such that $\mu(w) < \mathbf{d}(w)(\llbracket clear \rrbracket)$ for any choice of μ .

that either p or $clear(p)$ is in the common ground, since Barker does not assume that the value of $\mathbf{d}(w)(\llbracket clear \rrbracket)$ is part of the common ground. By asserting clarity, a speaker sets an upper bound on $\mathbf{d}(w)(\llbracket clear \rrbracket)$ and ensures that p and $clear(p)$ enter the common ground.

Bronnikov (2008) offers several arguments against Barker’s core proposal that the purpose of clarity assertions is to set the value of $\mathbf{d}(w)(\llbracket clear \rrbracket)$. First, Bronnikov notes that it is felicitous to assert the clarity of a necessarily true proposition (19).

(19) Take an integer n that is divisible by 9. It is clear that n is divisible by 3.

Any standard of clarity is consistent with it being clear that n is divisible by 3 if n is divisible by 9, so asserting the clarity of this proposition cannot tell us anything about $\mathbf{d}(w)(\llbracket clear \rrbracket)$. Thus, Barker’s proposal leaves us with no explanation for the function of an assertion of (19).

Second, Bronnikov points out that Barker’s proposal has trouble handling multiple clarity assertions in which the conclusions follow from similar bodies of evidence. Suppose that Bo and Jack see a woman walk by wearing a white coat and a stethoscope. Bo may felicitously utter (20a). Next, a man walks by dressed the exact same way. Jack may felicitously utter (20b).

- (20) a. BO: It’s clear that she’s a doctor.
 b. JACK: It’s clear that he’s a doctor too.

Since our evidence for concluding that the woman and man are doctors is highly similar, all the worlds in which the woman is a doctor should be just as normal as the worlds in which the man is a doctor. According to Barker, once Bo utters (20a), she sets the standard for clarity such that all worlds in which the woman is a doctor are at least as normal as this standard. Given this, once we see the man wearing a white coat and a stethoscope, it should immediately follow that all worlds in which the man is a doctor are at least as normal as the standard. But then (20b) should have no effects on the contextual standard for clarity, and thus no discourse effects whatsoever. Nonetheless, (20b) is felicitous.

2.3 Wolf and Cohen (2011)

Wolf and Cohen (2011) offer an analysis of clarity that is reminiscent of Barker and Taranto’s original analysis in that it is based on individuals’ beliefs.¹² Wolf and Cohen have two main reasons for adopting a belief-based analysis of clarity. First, examples such as (11) suggest that personal clarity statements reflect the beliefs of the experiencers. Second, Wolf and Cohen argue that *clearly* patterns with a class of adverbs described in Piñon (2006) that signal the strength of a speaker’s belief in the proposition that they modify. If we assume that *clearly* also modifies the strength of a speaker’s belief and that *clearly* and *clear* are closely related, then *clear* too should have something to do with belief. At the same time, Wolf and Cohen also follow Barker (2009) in assuming that mutual belief of all discourse participants is neither a necessary nor sufficient condition for simple clarity to be true. As a result, a statement of simple clarity cannot be reduced to a statement about the beliefs of the speaker or of discourse participants.

Wolf and Cohen reconcile the arguments for and against incorporating belief into the semantics of clarity by “objectivizing belief.” For personal clarity statements, only the beliefs of the explicit experiencer are relevant for determining clarity. When an experiencer argument is implicit, the relevant beliefs are not those of any particular discourse participants, but are rather the beliefs of “good reasoners” or experts. Formally, let R be a set of reasoners. Each reasoner $r \in R$ is

¹²The approach to clarity assertions advocated in Wolf (2014) is largely similar to that presented in Wolf and Cohen (2011). Due to these similarities, I do not discuss Wolf (2014) separately.

associated with a weight v_r such that $v_r \geq v_{r'}$ if and only if r is at least as good a reasoner as r' .¹³ For each $r \in R$, $P_r(p)$ is the subjective probability r assigns to the proposition p .¹⁴ We then define $P_{\text{justification}}$ as follows:

$$P_{\text{justification}}(p) = \sum_{r \in R} v_r P_r(p)$$

That is, a proposition’s “justification probability” is a weighted sum of the subjective probabilities assigned to this proposition by all reasoners. Again let $\mathbf{d}(c)(\llbracket \text{clear} \rrbracket)$ be a function from contexts c to degrees of likelihood. We have:¹⁵

- (21) a. *It is clear that p* is true in a context c iff $P_{\text{justification}}(p) \geq \mathbf{d}(c)(\llbracket \text{clear} \rrbracket)$.
b. *It is clear to x that p* is true in a context c iff $P_x(p) \geq \mathbf{d}(c)(\llbracket \text{clear} \rrbracket)$.

That is, a personal clarity statement is true if and only if the experiencer’s degree of belief in the conclusion is greater than this standard. A simple clarity statement is true if and only if the weighted sum of subjective probabilities of the conclusion over all individual reasoners is greater than the contextual standard for clarity.

Barker (2011) summarizes Wolf and Cohen’s proposal as one according to which “asserting clarity entails facts about what an expert would believe.” (192). We can view this proposal as a rejection of each of the assumptions in (5). Rather than assuming that if an assertion of *clear*(p) is felicitous, any agent can conclude p on the basis of evidence in the common ground, Wolf and Cohen claim that an assertion of *clear*(p) is felicitous so long as *good reasoners* would conclude p . Because Wolf and Cohen believe that agents can be weighted in terms of how good they are at reasoning, it follows that not all agents are fully rational.

Wolf and Cohen’s proposal suffers from several shortcomings. First, they provide no motivation for why simple clarity makes reference to weighted beliefs but personal clarity does not. In addition, one of Bronnikov’s critiques of Barker’s proposal also applies to Wolf and Cohen’s. On Wolf and Cohen’s account, when Bo asserts in (20a) that it is clear that a woman wearing a white coat and stethoscope is a doctor, she informs Jack that good reasoners would conclude that the woman is a doctor. Jack then asserts in (20b) that it is clear that a man wearing a white coat and a stethoscope is a doctor. Jack’s assertion should only be informative if there is a possibility that good reasoners would conclude that the woman is a doctor, but not the man. There is nothing logically inconsistent about good reasoners having such beliefs, but it would require us to make rather odd assumptions about the beliefs of good reasoners. More plausibly, if Bo’s assertion tells us that good reasoners would conclude that the woman is a doctor, we should already know that good reasoners would conclude that the man is a doctor. But this leaves us without an explanation for the function of Jack’s assertion.

Wolf and Cohen’s proposed solution to the paradox of asserting clarity also suffers from a more general problem. Namely, it is not obvious why we should care what experts would conclude, particularly if we are discussing conclusions that the discourse participants should be able to draw on their own. When Bo and Jack encounter the line outside Tartine, why does it matter that

¹³Wolf and Cohen suggest that R should include *all* reasoners in the domain of individuals, but this ignores the fact that most reasoners will not have the relevant evidence to assess the preagent of a clarity statement. For example, when Bo and Jack look at the line outside Tartine, most reasoners will not have access to this knowledge. To make sense of such cases, we must either restrict the set of reasoners R to those reasoners who possess the relevant evidence or have a reasoner r ’s reasoning weight v_r to depend upon whether r has access to the relevant evidence.

¹⁴Although Wolf and Cohen make use of subjective probabilities, the proposal would be essentially the same if it made reference to non-probabilistic credences or degrees of belief instead.

¹⁵As with the proposal of Barker and Taranto (2003) and Barker (2009), Wolf and Cohen’s proposal falls short by failing to incorporate the evidence supporting a clarity statement into the truth conditions of *clear*.

experts can conclude that there is a wait at Tartine? Surely Jack should be able to conclude this on his own, so it is still mysterious what role Bo’s clarity assertion is playing.¹⁶

2.4 This Missing Inference Hypothesis & Bronnikov (2008)

Although Barker and Taranto (2003) and Barker (2009) opt for the approaches to clarity assertions described above, they also consider an alternative option: the “missing entailment hypothesis” or “missing inference hypothesis.” Barker (2009) characterizes this proposal as follows:

Perhaps an assertion of clarity merely calls to the attention of the addressee some important fact already entailed by the common ground, but which has somehow not yet been added to the common ground. The evidence is available, the conclusion follows, but for some reason the addressee is hesitating to make that last step to the final conclusion. Perhaps they have a logic deficit, and can’t compute the consequences of their own beliefs (257–258).

A missing inference approach to clarity assertions can be seen as a rejection of (5b) and (5c), the assumptions that all discourse participants are rational and logically omniscient and that it is common ground that all discourse participants assume other participants to be rational and logically omniscient. On this view, asserting *clear*(p) when p can be inferred from the common ground addresses deficiencies in discourse participants’ rationality.

Barker and Taranto (2003) argue against this approach due to its incompatibility with the standard assumption that the common ground is closed under logical entailment. But this objection only works against a missing *entailment* hypothesis of clarity, as opposed to a missing *inference* hypothesis. Even on the assumption that the common ground is closed under entailment, it is possible that there is some proposition p that is not entailed by the common ground, but which a rational agent would plausibly infer from the common ground. Barker (2009) offers another criticism of the missing inference hypothesis, by noting that it would resemble von Fintel and Gillies’s (2011) proposal about epistemic *must*. Barker points out that *clear* and *must* differ in at least two respects. First, Barker claims that assertions of clarity depend on publicly available evidence, whereas assertions including epistemic *must* can be based on private evidence. Second, *clear* is gradable, whereas *must* is not. Whether or not these empirical observations are correct,¹⁷ it is hard to see how these differences are relevant to whether a missing inference hypothesis can be pursued for *clear*.

Bronnikov (2008) articulates a version of the missing inference hypothesis, proposing that *It is clear to x from e that p* is true if and only if x has performed a sound inference with e as premises and p as conclusion. In addition, this inference must be neither too difficult nor trivial. The prohibition against inferences that are too difficult stems from the intuition that if some inference is sound, but requires a great deal of mental effort to perform, its conclusion is not clear. The requirement that inferences are non-trivial is meant to capture the fact that clarity assertions are infelicitous if the prejacent has been explicitly stated in the preceding linguistic discourse:

- (22) a. Mindy is a doctor.

¹⁶Moreover, if we make the adjustment discussed in fn. 13 to restrict the relevant set of reasoners to those in possession of the relevant evidence, clarity statements may not even reflect the views of experts. If Bo and Jack are the only reasoners in possession of the relevant evidence, a simple clarity statement is simply about what Bo and Jack believe. In this case, it is even more mysterious what role the clarity assertion is playing.

¹⁷We will see below that it is possible to assert simple clarity on the basis of private evidence. Although epistemic *must* is non-gradable in standard theories of modality following Kratzer (1991), several contemporary theories of modality take modals to be highly similar to gradable adjectives (Lassiter 2016; Portner and Rubinstein 2016).

b. # Therefore, it is clear that Mindy is a doctor.

Bronnikov formalizes this proposal using a dynamic logic for inferences adapted from Duc (2001). $\mathbf{B}_x p$ is true if and only if x believes p . For a set of agents X , $\mathbf{B}_X p$ is true if and only if $\forall x \in X (\mathbf{B}_x p)$. Note that the contents of agents' beliefs are syntactic formulas, and there is no assumption that agents' beliefs are closed under logical entailment or rational inference. Thus, even if p entails q and $\mathbf{B}_x p$ is true, it does not follow that $\mathbf{B}_x q$. The formula $\langle a \rangle \varphi$ is true if and only if φ holds after the action a is performed. This can be generalized for sets of actions: $\langle A \rangle \varphi$ means that φ holds after some subset of actions $a \subseteq A$ are performed.

Now let \mathcal{A} represent the set of all actions and let X represent the set of discourse participants in some context. Inferences are a subset of actions and are partitioned into three sets depending on their difficulty relative to some particular agent. Let $Trivial_x \subseteq \mathcal{A}$, $Easy_x \subseteq \mathcal{A}$, and $Difficult_x \subseteq \mathcal{A}$ represent the sets of trivial, easy, and difficult inferences, respectively, performed by x . We then have the following:

(23) *It is clear to x that p* presupposes $\neg \langle Trivial_x \rangle \mathbf{B}_x p$ and is true iff $\langle Easy_x \rangle \mathbf{B}_x p$.

(24) *It is clear that p* presupposes $\neg \langle Trivial_X \rangle \mathbf{B}_X p$ and is true iff $\langle Easy_X \rangle \mathbf{B}_X p$.

On this view, a speaker asserting simple clarity indicates that an inference is available to the other discourse participants. Again, since agents' beliefs are not closed under entailment or rational inference, addressees may possess evidence sufficient to draw some inference yet fail to do so.

Bronnikov argues that this proposal does not succumb to the two phenomena that proved problematic for Barker (2009): assertions of the clarity of necessarily true propositions (19) and sequential clarity assertions based on similar bodies of evidence (20). Since agents' beliefs are syntactic objects and are not closed under logical entailment, an agent may fail to believe a necessarily true proposition. If a speaker believes that an addressee does not believe a necessarily true proposition, the speaker may use a clarity assertion to indicate that this proposition may be soundly inferred. In the case of sequential clarity assertions, even if two propositions p and q are based on similar bodies of evidence, there is no guarantee that an agent who believes p will believe q . Each of p and q must be inferred independently in order for an agent to believe both. Thus, assertions of the clarity of p and q indicate that two distinct inferences are available to interlocutors.

Despite Bronnikov's success in handling these issues, Bronnikov's proposed truth conditions face several problems. First, Bronnikov informally claims that *It is clear to x from e that p* is true if and only if x has performed a sound inference with e as premises and p as conclusion. Although this characterization takes the evidential sufficiency requirement seriously, Bronnikov's formal truth conditions make no reference to the evidence argument of *clear* and therefore do not satisfy the evidential sufficiency requirement.

Second, Bronnikov adopts the position that when the value of x is left implicit, it defaults to the set of all discourse participants. It follows that if we let X be the set of all discourse participants and assume that *It is clear that p* is true, we have $\langle Easy_X \rangle \mathbf{B}_X p$. Bronnikov writes that this formula should be interpreted to mean that every individual in X has performed an easy inference after which they all believe p . But if this is the correct interpretation of $\langle Easy_X \rangle \mathbf{B}_X p$, then any context in which *clear*(p) is truthfully asserted, every discourse participants must have already concluded p before *clear*(p) was asserted. This result undermines Bronnikov's claim that clarity assertions highlight an inference that an addressee has failed to infer and seems to recreate the paradox of asserting clarity.

One approach to resolving this issue would involve giving a different interpretation to $\langle A \rangle \varphi$. Assume now that $\langle Easy_X \rangle \mathbf{B}_X p$ means that for every individual $x \in X$, there exists an easy inference for x which, if performed, would result in x believing p . Perhaps only the speaker of a simple

clarity assertion has *actually* performed the relevant inference, but the inference is available to all participants. This reinterpretation of $\langle A \rangle \varphi$ means that a speaker can assert the clarity of p truthfully in a context in which an inference whose conclusion with p has not been performed by all discourse participants, but is available to all discourse participants. Unfortunately, this interpretation of $\langle A \rangle \varphi$ generates a new problem for personal clarity claims. Suppose a speaker asserts *It is clear to me that p* , which on Bronnikov’s account presupposes that $\neg \langle Trivial_{speaker} \rangle B_{speaker} p$. On the current interpretation of $\langle A \rangle \varphi$, this presupposition means that there is no trivial inference available to the speaker after the performance of which the speaker would believe p . But presumably a speaker who asserts *It is clear to me that p* already believes p , in which case there is a trivial inference available to the speaker that takes p as the premise and p as the conclusion. It seems that there is no way to give a consistent interpretation to $\langle A \rangle \varphi$ that gives the desired result for both simple and personal clarity.

Perhaps the above argument simply suggests that it is incorrect to take the lexical semantics of *clear* to presuppose that the inference involved is not trivial. There are additional reasons to question this presupposition. If I look outside the window and see that it is pouring down rain, I can truthfully assert that it is clear that it is raining. If any inference is involved in this case, it is the relatively trivial inference from my observation that water is falling from the sky to the conclusion that it is raining. It may be the case that to assert that it is clear that it is raining would be pragmatically odd in many contexts in which the speaker and interlocutor have direct evidence of the rain. Nonetheless, this oddness should not be explained via the truth conditions of *clear*.¹⁸

In addition to these concerns, Bronnikov’s emphasis on inference rules is problematic for cases in which it seems inappropriate to say that inference is involved at the time clarity is asserted. For example, there are cases in which the speaker and addressee reached the conclusion of a clarity assertion well before clarity was asserted. Consider the following example from the Corpus of Contemporary American English (COCA; Davies 2008):

- (26) It is clear that we have a highly polarized, very sharp, and I think at times, too sharp level of anger expressed across the political aisles in this country.

Example (26) was spoken by journalist Juan Williams on Sean Hannity’s television program in 2011. Presumably, both Williams and Hannity had already come to the conclusion that there was a highly polarized level of anger in American politics. On Bronnikov’s account, Williams was signaling to Hannity that Williams had performed some type of inference or that such an inference was available. This seems to miss the point of Williams’s assertion, which was more plausibly an attempt to draw Hannity’s attention to something Hannity *already* believed.

Still, Bronnikov’s approach successfully handles cases that are left unexplained by Barker (2009) and Wolf and Cohen (2011). The final answer to the paradox of asserting clarity should preserve this feature while at the same time addressing the shortcomings of Bronnikov’s proposal.

¹⁸An earlier version of this article used an example in which the prejacent of *clear* contained *might*. This example was intended to illustrate that a very weak conclusion can appear within the scope of *clear* and that any inference to such a weak conclusion must be relatively trivial. The example was removed to streamline this discussion, but I note that an anonymous reviewer who did not find the original example particularly felicitous suggests that there may be a prohibition on *might* appearing within the scope of *clear*. Contra this claim, attested examples of *might* appearing within the scope of *clear*, such as the following, can be found:

- (25) ... it is clear that another man might find Kepler’s theory, that the celestial spheres are proportional to the inscribed and circumscribed spheres of the different regular solids, more agreeable to *his* reason (Peirce 1877.)

3 A New Semantics for Clarity

While none of the proposals discussed above are fully satisfactory, they do provide useful insights that will assist in developing a new semantics for *clear*. I follow Barker and Taranto (2003), Bronnikov (2008), and Wolf and Cohen (2011) in taking clarity to be essentially tied to the beliefs of the experiencer. If the experiencer argument includes more than one individual, then the prejacent must be clear to all members of the experiencer. In order to account for the vagueness of *clear*, I follow Barker and Taranto (2003), Barker (2009), and Wolf and Cohen (2011) in taking the truth of a clarity statement to depend on some contextually relevant threshold $\mathbf{d}(c)(\llbracket clear \rrbracket)$. In addition, I incorporate a version of the evidential sufficiency requirement. As discussed above, this states that the evidential argument of *clear* must provide “sufficient” grounds for the experiencer to have a suitably high degree of belief in the prejacent. Exactly what counts as “sufficient” in this sense remains to be spelled out in detail. I return to this question below. Combing these ideas, I offer the following first pass, relatively informal proposal regarding the meaning of *clear*:

- (27) Truth conditions for personal-evidential clarity (first version, to be revised):
It is clear to x from e that p is true at a world w and a context c iff for all $i \in x$, i possesses the evidence e at w , i 's degree of belief in p at w is at least $\mathbf{d}(c)(\llbracket clear \rrbracket)$, and i 's possession of e is sufficient grounds for i 's degree of belief in p to be at least $\mathbf{d}(c)(\llbracket clear \rrbracket)$.

The truth conditions for personal, evidential, and simple clarity statements can be derived from those in (27) by establishing how the evidence and experiencer arguments of *clear* are interpreted when left implicit. Negated clarity statements suggest that implicit evidence arguments are interpreted existentially:

- (28) It is not clear to me that Mindy is a doctor.

Example (28) entails that there is no evidence from which it is clear to the speaker that Mindy is a doctor. This is consistent with taking implicit evidence arguments to receive an existential interpretation that scopes under negation.¹⁹

Assuming that implicit evidence arguments receive an existential interpretation, we face the following question. If the experiencer argument of *clear* includes more than one individual and the evidence argument is left implicit, is it the case that the prejacent is clear to each member of the experiencer on the basis of the same body of evidence or is it possible that these experiencers' bodies of evidence differ? This question can be seen as one related to the scope ordering of two quantifiers. Suppose x is a set of individuals, and let $clear(i, e, p)$ denote that p is clear to i from e . Then, we have two candidates for the truth conditions of *It is clear to x that p*:

- (30) a. $\forall i \in x (\exists e (clear(i, e, p)))$
 b. $\exists e (\forall i \in x (clear(i, e, p)))$

Bronnikov (2008) argues for (30b) on the basis of the following example. Suppose three individuals have each read a different novel by Dostoyevsky, and all three have concluded that Dostoyevsky

¹⁹It is not uncommon for implicit arguments to be interpreted existentially. Condoravdi and Gawron (1996) provide the following example:

- (29) There was a piece of bread on the table, but John didn't eat.

It is a priori plausible that the implicit argument of *eat* would be interpreted as the contextually salient piece of bread. If this were the case, then (29) would be true in a situation in which John did not eat the piece of bread, but ate something else. However, (29) is not compatible with such a situation. To arrive the correct interpretation for (29), the implicit object of *eat* must be interpreted existentially.

is a genius. Bronnikov claims that it is not the case that it is clear to all three that Dostoyevsky is a genius. Rather, according to Bronnikov, for it to be clear to all three that Dostoyevsky is a genius, all three must share the *same* body of evidence for reaching this conclusion. In contrast, my own judgments are that in the situation described by Bronnikov, it *is* clear to all three that Dostoyevsky is a genius. Thus, I assume the scope ordering shown in (30a).

3.1 Implicit Experiencers & Predicates of Personal Taste

While the interpretation of implicit evidence arguments of *clear* is relatively straightforward, the interpretation of implicit experiencer arguments poses a greater challenge. Most authors have taken the implicit experiencer argument to be interpreted as including all discourse participants, but Barker (2009) points out that this is not always the case. Suppose Bo shows up to a faculty meeting drunk. Bo knows that Jack knows that she is drunk, but is unsure as to whether the other meeting attendees knows this. Bo can turn to Jack and ask the following:

(31) BO: Is it clear that I'm drunk?

Bo is not asking whether it is clear to the discourse participants, which including only herself and Jack, that she is drunk. Rather, she seems to be asking whether it is clear to everyone else in the meeting that she is drunk, even though these individuals are not participants in the discourse. A similar point can be made with past tense clarity statements, such as the following web examples:

(32) Even in the nineteenth century it was clear that racial discrimination in jury selection affected the due process rights of African Americans.

(33) Six months later, Hitler invaded Czechoslovakia, in outright defiance of the Munich agreement. It was clear that Hitler's next target was Poland, and Britain and France pledged themselves to its defense.

The statements in (32) and (33) do not entail that the prejacent were clear to the authors or their readers in the nineteenth century or 1936, respectively. Rather, these examples are best understood as claiming something about what was clear to a certain set of experiencers at the time.

One way to account for the existence of these “exocentric” uses of *clear* is to assume that an implicit experiencer argument takes some contextually-determined value that is often, but not always, a set of individuals that includes all discourse participants. However, this view is challenged by the following situation suggested by an anonymous reviewer. Suppose one hundred people are in a room looking at a photograph of a woman. Ninety-nine of the individuals conclude on the basis of the photograph that it is clear that she is a doctor, while there is one dissenter. If we took the implicit argument of *clear* to default to some contextually relevant set of individuals, then it would be plausible to assume that all individuals in the room would be members of this contextually relevant set. Given the one dissenter, it would not be clear to all individuals in the room that the woman is a doctor. However, the reviewer and I share the judgment that in this situation, it may be felicitous to assert that it is clear she is a doctor.²⁰ This example may be seen as evidence for Wolf and Cohen's objectivized belief model, according to which leaving an experiencer argument implicit forces us to consider the weighted beliefs of reasoners. On this view, the views of the majority in the room can outweigh those of the lone dissenter. However, as noted above, this analysis offers no

²⁰A reviewer points out that such a situation might license one to say that it is clear that the woman is *not* a doctor. In particular, if the view of the one dissenter is privileged by the context, then the dissenter's perspective may be the only one that is relevant. See the discussion of example (40) below for a case in which the viewpoints of some perspectives are privileged.

explanation for why experiencers' reasoning abilities are weighted when the experiencer argument is left implicit, but not when it is explicit.

Both the objective belief analysis and the contextualist analysis run into problems accounting for examples in which it appears that counterfactual situations are relevant for determining the clarity of propositions. Consider the following scenario, based on an example from Kripke (1980). Suppose Bo uncovers documents that lead her to believe that a mathematician named Schmidt was the true author of Gödel's famous incompleteness theorems. Gödel stole the proofs from Schmidt and passed them off as his own. Bo is the only one who is in possession of these documents when she utters the following:

(34) BO: It is clear from these documents that Gödel was a fraud.

Since the only individual in possession of the relevant evidence is Bo, on the contextualist account the only plausible value for the implicit experiencer argument is Bo. On the objectivized belief account, to make sense of this example we must either restrict the set of reasoners to include only Bo or give very little weight to the beliefs of reasoners who are not in possession of the relevant documents.²¹ On either view, (34) should be true if and only if (35) is true:

(35) BO: It is clear to me from these documents that Gödel was a fraud.

In the context described, (35) is certainly true. Despite this, we may still have doubt about the truth-value of (34). For example, it may be that Bo has misinterpreted the documents and that others would not reach the same conclusion as her. To evaluate the truth of (34), it seems necessary to consider what other agents *would* conclude if they were in possession of the relevant evidence. In other words, the truth conditions of clarity statements with implicit experiencers appear to depend upon counterfactual situations.

I propose that we can make sense of the above phenomena by first noting the similarities between *clear* and predicates of personal taste (PPTs). The idea that *clear* and PPTs may be related is not novel; both Barker (2009) and Wolf and Cohen (2011) suggest that clarity statements may be related to PPTs, although they stop short of claiming that *clear* is a PPT. Like *clear*, statements involving PPTs may or may not be relativized to a particular experiencer:

(36) This cake is tasty (to me).

PPTs and *clear* are similar in other respects as well. As noted above, a proposition *p* may be clear due to the judgments of a large majority of relevant individuals, even if there remain dissenters who hold that *p* is not clear. Similarly, we find that a PPT may be applicable in a context in which most experiencers agree that it applies, notwithstanding some dissent. Consider the following example, modified from Keshet (2005):

(37) BO: The chili is tasty.

JACK: No, the chili is not tasty.

BO: Wait, how can you say that? The beans are bursting with flavor and the meat is cooked to perfection!

JACK: Well, it's too spicy, for one thing.

BO: Let's ask someone else. . .

Keshet points out that because the discourse in (37) is coherent and rational, it makes sense to base one's judgment on the applicability of a PPT on the opinions of people in general, despite the existence of some dissenters.

²¹See fn. 13 for reasons why Wolf and Cohen's proposal may in general require such adjustments.

PPTs and *clear* are also alike in that they both exhibit exocentric uses. In most situations, a speaker’s assertion that chili is tasty implies that the chili is tasty to the speaker. Yet there are other cases in which a speaker may assert that a PPT holds of some entity without implying that the PPT relativized to the speaker holds of that entity. Consider the following example, based on an example from Stephenson (2007):²²

(38) BO: This brand of cat food is tasty. Whenever I feed it to Spot, she eats the whole bowl.

In (38), Bo is not committed to finding the cat food tasty to herself. Rather, she asserts something more or less equivalent to the proposition that the cat food is tasty to cats.

Of course, simply pointing out the similarities between *clear* and PPTs does not resolve the issue of how to account for implicit experiencer arguments of *clear*. What is needed is some theory that accounts for the above PPT data that can be extended to *clear*. While there is no consensus regarding the correct semantic analysis of PPTs, here I draw on recent genericity-based analyses of PPTs in addressing the issue of how to resolve implicit experiencer arguments of *clear* (Keshet 2005; Anand 2008; Moltmann 2010; Pearson 2013). To motivate this choice, I demonstrate below how a genericity-based account is well-suited to handle many of the issues that arise when determining how to resolve implicit experiencer arguments. While I take inspiration from genericity-based analyses of PPTs, I do not directly adopt any previous proposal, although the proposal below is particularly indebted to Pearson (2013)’s argument that PPTs involve genericity and quantification over both worlds and individuals.

Let $clear(x, e, p, w, c)$ be true if and only if the proposition that p is clear to x from e at world w in context c , following (27). Let Acc be an accessibility relation between worlds, and let $Rel(w, c)$ be a function from worlds w and contexts c to the set of relevant individuals in w and c . Finally, let GEN be a generic quantifier (Krifka et al. 1995). Then, a clarity statement in which the experiencer is left implicit is interpreted as follows:

(39) Truth conditions for evidential clarity:
It is clear from e that p is true at a world w and a context c iff
 $GENw' GENx (Acc(w, w') \wedge x \in Rel(w', c) \rightarrow clear(x, e, p, w', c))$

In other words, a proposition p is clear on the basis of some evidence e if and only if it would (in general) be clear (in general) to relevant experiencers. As discussed above, if the the evidence argument is also left implicit, it receives a narrow scope existential interpretation.

First, a few words on who counts as a relevant individual. In the case of PPTs, Pearson (2013) points out that, at a minimum, relevant individuals must have direct experience with the entity of which the PPT is predicated. For example, in evaluating the tastiness of a particular type of chili, only individuals who have tasted the chili will be relevant. But there are other constraints on who counts as a relevant individual. In most cases, only cats are relevant experiencers for judging the tastiness of cat food, even if Bo may have once accidentally tasted her cat’s food.

We can observe similar constraints in the case of *clear*. As discussed above, relevant experiencers only include individuals in possession of the relevant evidence. For example, when Bo and Jack discuss whether it is clear that there is a wait at Tartine, the relevant set of individuals should include only those who know that there is a long line outside Tartine. Moreover, just as the context may dictate whether it is the tastes of humans or the tastes of cats are relevant for determining whether some food is *tasty*, context may dictate exactly what types of reasoners’ beliefs are relevant for determining whether some proposition is clear. Barker (2011) provides an example in which two chess grandmasters are watching a game when one utters the following:

²²Stephenson’s original example was suggested by Kai von Fintel (p.c.).

(40) It is clear that white will win in nine moves.

Most individuals, even those who know the rules of chess, are probably not able to conclude that white will win in nine moves. But given a context in which both interlocutors are grandmasters, we can take the relevant individuals to only include those with expert knowledge of chess.

As alluded to above, the proposal in (39) is able to explain several outstanding questions regarding clarity statements. Exocentric uses arise when the contextually relevant set of individuals includes non-discourse participants. Note that even if the relevant evidence is only possessed by discourse participants in the actual world, the definition in (39) requires that we consider what relevant, non-discourse participants would conclude in accessible worlds in which they do possess the evidence. This explains why in the case of (34), the truth of Bo's statement depends on whether it is true in general that other experiencers would conclude that Gödel was a fraud in those worlds in which they are relevant in virtue of possessing the evidence.

The proposal in (39) also explains why it may be felicitous to assert the clarity of p , even when p is not clear to some relevant individuals. Because generics allow for exceptions (Krifka et al. (1995)), p may be clear to the relevant experiencers *in general* even if it is not clear to *all* relevant experiencers. In fact, it may be the case that the speaker herself does not believe the prejacent of her own clarity assertion. Thus, the current proposal is consistent with Barker's (2009) claim that a speaker may assert (41), repeated from (17), while still privately believing that there is life on Mars.

(41) It is clear that Mars is barren of life.

Still, a speaker who does believe there is life on Mars can only assert (41) if she is prepared to acknowledge her beliefs diverge greatly from those of most relevant people. In his discussion of example (41), Barker also notes that simple clarity statements have a normative quality, in that they imply that p is what is rational to conclude. The proposal in (39) may offer an explanation of this quality of clarity statements. Assume that whenever individuals in general would conclude that p is likely on the basis of some evidence, it is rational to conclude p on the basis of this evidence. Then, given the proposal in (39), whenever p is clear, it is rational to conclude p .

The irrationality of a speaker who asserts the clarity of p but does personally believe that p is clear offers one explanation of why simple assertions of clarity imply that the speaker believes the prejacent. Another explanation for this implication comes from consideration of what licenses a speaker to make a simple clarity claim, i.e. what licenses an individual to make a claim about what is clear to people in general. Here, another comparison with PPTs is instructive. López de Sa (2008) discusses a principle known as the "Presupposition of Commonality," which states that individuals are roughly similar in their tastes. Pearson (2013) invokes this principle to explain how speakers are able to make simple claims involving PPTs (e.g. *This chili is tasty*) when, on her analysis, such claims are about the tastes of people in general. Usually, speakers are licensed to make such claims because they have direct knowledge of their own tastes and can generalize to the tastes of others by assuming the Presupposition of Commonality. In other words, because I know this chili is tasty to me and I know that my tastes are similar to those of others, I can claim that this chili is tasty to people in general. Then, when a speaker asserts that the chili is tasty simpliciter, listeners can reason that the speaker's claim about what is tasty to people in general was justified on the basis of the speaker's own tastes. Extending this type of reasoning to the case of clarity assertions, we assume that individuals are roughly similar in their reasoning and inferential capacities. In many cases, this assumption allows a speaker to reason from the fact that a proposition p is clear to the speaker to the stronger claim that p is clear simpliciter. Listeners can then reason that if a speaker asserts the (simple) clarity of p , it is likely the case that p is clear to the speaker.

Before moving on, I note that one of the hallmark features of PPTs is that they lead to cases of faultless disagreement (Lasersohn 2005). Consider the first two lines of (37). Bo and Jack seem to have a genuine disagreement about whether the chili is tasty, but supposing they have each spoken genuinely, it does not appear that either has said something false. Given the similarities between *clear* and PPTs, we would expect *clear* to also be involved in cases of faultless disagreement. While there are genuine cases of faultless disagreement involving *clear*, in some cases it may be more difficult to detect faultlessness in disagreements over clarity claims that it is in disagreements involving PPTs. Suppose Bo and Jack see a woman wearing a white coat and smoking a cigarette, at which point they have the following exchange:²³

- (42) BO: It's clear that she's a doctor, since she's wearing a white coat.
 JACK: No, it's not clear that she's a doctor, since she's smoking.

One may have the sense that in the disagreement between Bo and Jack shown in (42), one of the two is, in fact, wrong. In other words, we may think that disagreement is not faultless.

To explain why some disagreements over clarity assertions may not appear to be faultless, first recognize that when individuals disagree about the clarity of p , they often disagree about the truth of p itself. If p does not contain a PPT, then a disagreement about the truth of p should not be faultless; one of the two speakers must be wrong. Now assume the version of the Presupposition of Commonality for *clear*, stating that individuals are similar in their inferential and reasoning capacities. It follows that if some relevant individual believes that p is false, then it is implausible that relevant individuals will in general assign p high enough degrees of belief such that p is clear to them.²⁴ To illustrate this line of reasoning with the example in (42), note that Jack's utterances suggests that he believes that the woman is not a doctor. Since Jack believes she is not a doctor and we assume that the reasoning capacities of people in general do not greatly differ from those of Jack, it is hard to imagine that people will in general find it clear that she is a doctor. Thus, we may have the sense that in this situation Bo is at fault. Indeed, this line of reasoning is discussed by Barker and Taranto (2003) and Barker (2009), who claim that if an interlocutor disagrees with a speaker's claim that p is clear, then p is not clear.

That said, modifying (42) provides a relatively straightforward case of faultless disagreement. Consider a situation in which Bo and Jack are looking at a woman dressed in a white coat and have the following exchange:

- (43) BO: It's clear that she's a doctor, since she's wearing a white coat.
 JACK: While I agree she's likely a doctor, it's not clear that she is.

Now Bo and Jack are in agreement that the woman is likely a doctor; their disagreement is simply over whether it is clear that she is. As a consequence, in (43), it appears that both Bo and Jack to have said something true, i.e. to have a faultless disagreement. Thus, the prediction that clarity statements should give rise to cases of faultless disagreement just as PPTs do is borne out.

3.2 Informative Clarity Assertions

Given the semantics for clarity statements proposed above, it follows that clarity statements may be both truth and informative, contra the paradox of asserting clarity outlined in §1. On the semantics proposed above, (5a) and (5a') are no longer valid. Recall that these assumptions claim that the existence of evidence e in the common ground such that any rational agent can conclude p from e is either at least a necessary, and perhaps also sufficient, condition for *clear*(p) to be

²³This example is based on a similar situation discussed by Wolf and Cohen (2011).

²⁴As we have seen, this can be overcome if p is clear to a large body of relevant individuals.

true. The new claim is roughly that an assertion of *clear(p)* is felicitous so long as there is some evidence *e* such that in general the relevant individuals would conclude *p* from *e*. This allows for cases in which not all relevant individuals conclude *p* from *e*, such as the scenario described above in which ninety-nine out of one hundred discourse participants agree it is clear that a woman in a photograph is a doctor. In these cases, *p* may not be mutually believed by all discourse participants even though *clear(p)* may be felicitously asserted. In such a case, the addition of *clear(p)* to the common ground may inform those dissenters about the beliefs of individuals in general and may even serve as an inducement for those hold-outs to change their views.

Simple clarity assertions can also be informative in situations in which not all discourse participants are relevant. In such cases, we have no reason to believe that discourse participants will already believe the prejacent or even have access to the relevant evidence that justifies belief in the prejacent. Consider the following attested examples from COCA:

- (44) a. It is clear that Maliki has come out as the winner in the political crisis he provoked. He has made it more difficult for his Shia rivals to dissent while simultaneously confining his Sunni opponents in a position suitable for exerting pressure and exploiting divisions within their ranks.
- b. And it is clear that even in flush times, the Olympics carry a considerable financial burden. The 1992 Barcelona Games left Spain with a \$6.1 billion debt. Athens estimated that the 2004 Games would cost \$1.6 billion, but in the end it was \$16 billion. Meanwhile, it took Montreal nearly 30 years – until 2005 – to pay off the \$2.7 billion it owed after the 1976 Summer Games.

The examples in (44) come from news contexts in which, intuitively, the relevant individuals are experts who have access to the relevant evidence. Importantly, the addressee(s) in these contexts cannot be assumed to already possess the evidence, since this evidence is provided immediately after clarity is asserted. If the addressee(s) possessed this evidence before clarity was asserted, the statements made following the clarity assertions would be completely superfluous. Note that in these cases, a clarity assertion updates the common ground information about the beliefs of the implicit experiencer, and possibly with information about the evidence supporting those beliefs. Depending on how trustworthy we take the experiencer to be, the prejacent may also enter the common ground, but this is not guaranteed.

While such examples demonstrate that even simple clarity assertions may be informative, the most interesting types of clarity assertions are those that are not obviously informative in the ways described above. We can find examples of clarity assertions in which all discourse participants are relevant, they all have access to the necessary evidence, and there is no reason to think that any of them would fail to conclude that the prejacent is the case. For example, consider our original motivating example in which Bo and Jack encounter a line outside Tartine when Bo asserts that it is clear there is a wait at Tartine. Both Bo and Jack can be assumed to be relevant experiencers, and the evidence supporting Bo's assertion is available to both of them. It seems reasonable to think that Jack can conclude that there is a wait at Tartine on his own, so what purpose does Bo's assertion serve? In the next section, I show how we can understand clarity assertions in such cases as being used to raise awareness of an issue that the speaker thinks the addressee may not be attending to.

4 Clarity & Awareness

To handle cases in which it does appear that all discourse participants are relevant, I argue that we should abandon the second two assumptions underlying the paradox of asserting clarity: (5b) and

(5c). Recall that these assumptions held that all discourse participants are rational and logically omniscient and that it is common ground that all discourse participants are rational and logically omniscient. By abandoning these assumptions, my approach resembles the missing inference hypothesis, a version of which was defended by Bronnikov (2008). However, it was noted above that there was no consistent way to interpret Bronnikov’s formal proposal and that Bronnikov’s proposal placed too much emphasis on explicit inferences. The approach outlined below relies upon a more general notion of agents’ awareness and attentional states and the possibility of these states being manipulated by utterances. The key idea of the proposal put forward below is that linguistic agents’ rational capacities are limited by their states of awareness. Because of this limitation, discourse participants may fail to believe p even when the evidence available to them is sufficient to conclude p . An agent’s failure to conclude p may stem from unawareness of the available evidence or from unawareness of the conclusions that follow from this evidence. In turn, agents’ unawareness can be overturned if another agent asserts the clarity of p .

Unawareness has been well-studied in the fields of artificial intelligence, economics, and rational choice theory, particularly following Fagin and Halpern (1987). Fagin and Halpern distinguish several different types of unawareness. One form of unawareness consists in an agent’s lack of conceptual grasp. If an agent lacks the concept of a croissant, the proposition *Tartine has great croissants* will not be comprehensible to the agent regardless of the agent’s state of attention. Unawareness may also involve processing limitations. For example, an agent might encounter a long and complex logical formula that uses only negation, conjunction, and the proposition *There is a line outside Tartine*. The agent may attend to this formula and understand the meaning of each of its components, but still remain unaware of what the full formula’s meaning is due to lack the resources to process it within a reasonable amount of time.

Unawareness due to inattention is the most relevant form of unawareness for my purposes and involves neither lack of conceptual grasp nor processing difficulties. Instead, unawareness due to inattention reflects the fact that at any given time agents are unable to attend to all issues which they possess the conceptual and processing resources to understand. A growing body of work has argued that unawareness due to inattention is relevant to a variety of semantic, pragmatic, and discursive phenomena. This work includes discussion of “conversational backoff” (Rawlins 2010), work on the effects of questions on agents’ behavior (Franke and de Jager 2011), discussion of implicit and explicit belief for understanding the semantics of epistemic modals (von Fintel and Gillies 2010; Yalcin 2011), proposals regarding the “attentive content” of utterances (Ciardelli et al. 2011; Roelofsen 2013), and “epistemic resistance moves” in discourse (Bledin and Rawlins 2016). Below, I draw heavily from the ideas and formalisms of de Jager (2009) and Franke and de Jager (2011) in order to illustrate how clarity assertions may affect addressees’ awareness. The choice to follow the formalism of de Jager (2009) and Franke and de Jager (2011) is motivated primarily by Franke and de Jager’s focus on multi-agent interactions and the consequences that awareness and unawareness of issues has for decision making.

4.1 Unawareness Due to Inattention

To illustrate the intuitions behind Franke and de Jager’s model of awareness, we can consider a situation described in Franke and de Jager (2011). Suppose Little Bo Peep has lost her keys and doesn’t know where to find them. She looks in her pockets, on her nightstand, and in the sofa cushions, but they don’t show up. She has given up the search when, from his corner, Little Jack Horner offers the following help:

(45) JACK: Did you leave them in the car?

Bo slaps her forehead and runs out to look for her keys in the car.

Franke and de Jager characterize the change in Bo’s mental state that occurs in response to Jack’s utterance as Bo becoming aware of a new possibility. Before Jack speaks up, Bo exhibits unawareness due to inattention, which is characterized by the following properties. First, Bo’s initial unawareness of the possibility that her keys are in the car is not a form of uncertainty. Even if she believed there was a slight chance that the keys were in the car, she would go check there once she had checked the more likely locations. Second, Bo lacks introspection about her unawareness state; she does not know that she is unaware of some relevant possibility. Third, her unawareness is easily overturned, as illustrated by the fact that as soon as Jack mentions the possibilities of the keys being in the car, Bo becomes aware of this possibility.²⁵ Finally, Bo’s awareness state affects her approach to resolving her decision problem. Before Jack’s utterance in (45), she will look in the locations that she is aware of: her pockets, her nightstand, the sofa, etc. But after Jack speaks up, her decision problem changes insofar as she now has another action available to her: she can look in the car.

This notion of (un)awareness allows us to distinguish EXPLICIT BELIEF from IMPLICIT BELIEF. In the scenario described above, Bo has an explicit belief that her keys are not on her nightstand because she is aware of the possibility that her keys are on her nightstand, but believes they are not. Bo also possess beliefs that involve some form of unawareness; these are her implicit beliefs. Some of her implicit beliefs are also accompanied with an ASSUMPTION, while others are not. When Bo searches for her keys, she lacks any explicit belief that her keys are not in her car, but she acts as if she held such a belief. For this reason, we say that Bo has an implicit assumption that her keys are not in the car. Importantly, implicit assumptions are beliefs that an agent would not hold if she were to consider them explicitly.²⁶ As soon as Bo becomes aware of this proposition, her implicit assumption is overturned; she no longer behaves as if she believed her keys were not in her car.

An agent can also have an implicit belief without making an implicit assumption. To borrow an example from Yalcin (2011), suppose that Bo does not live in Topeka. Because she does not live there, it’s unlikely that Bo has any reason to consider whether or not it is raining in Topeka. But unlike the case with the keys, Bo’s behavior does not reflect any implicit assumption about whether it is raining in Topeka. We can say that Bo has an implicit belief that it might be raining in Topeka, but that this implicit belief comes with no implicit assumption. If someone were to ask Bo whether she thinks it is raining in Topeka, she would come to have an explicit belief about the issue. But it is very likely that her explicit belief would match her implicit belief, namely that it might be raining in Topeka.

To formalize these notions, Franke and de Jager first model an agent’s doxastic state under full awareness, i.e. the agent’s doxastic state were the agent fully aware of every proposition relevant to the current discourse. This full model is then filtered through an agent’s awareness state to yield the agent’s doxastic state under unawareness. Let \mathcal{W} be a set of possible worlds and let $\mathcal{P} \subseteq \wp(\mathcal{W})$ be a set of propositions.²⁷ An agent’s background model is defined as follows:

²⁵Note that the stimulus which produces a change in awareness need not be linguistic. Bo could look out her window, see her car, and realize that the keys might be in the car.

²⁶Here it is necessary to distinguish between implicit beliefs and implicit assumptions. As I discuss below, agents would retain many of their implicit *beliefs* if they were to consider them explicitly. But implicit *assumptions* are specifically those implicit beliefs revealed through agents’ actions that they would revise upon explicit consideration.

²⁷The model presented here is a “semanticized” version of the model used in Franke and de Jager (2011). In Franke and de Jager (2011), P is a set of atomic proposition letters. The semantics for these proposition letters are delivered via a valuation function $v : \mathcal{W} \times \mathcal{P} \rightarrow \{0, 1\}$. The move to treating propositions as sets of possible worlds is intended simplify the model, and does not reflect any deep theoretical differences from Franke and de Jager’s proposal.

- (46) An agent's BACKGROUND MODEL \mathcal{M} is a quadruple $\langle \mathcal{W}, \mathcal{A}, P, \mathcal{U} \rangle$ where:
- \mathcal{W} is a set of worlds.
 - \mathcal{A} is a set of actions.
 - P is a subjective probability distribution over \mathcal{P} .
 - $\mathcal{U} : \mathcal{W} \times \mathcal{A} \rightarrow \mathbb{R}$ is a utility function.

Next, we define an agent's awareness state, which will define those propositions and actions that an agent is unaware of, along with any implicit assumptions.

- (47) An agent's AWARENESS STATE α is a triple $\langle \mathcal{U}, \mathfrak{A}, \mathfrak{V} \rangle$ where:
- $\mathcal{U} \subseteq \mathcal{P}$ is a set of UNMENTIONABLE propositions.
 - $\mathfrak{A} \subseteq \mathcal{A}$ is a set of UNPERFORMABLE actions.
 - $\mathfrak{V} \subseteq \mathcal{U}$ is a set of ASSUMPTIONS.

We require that \mathcal{U} is closed under complement, i.e. if there exist two propositions $p, q \in \mathcal{P}$ such that $p = \mathcal{W} \setminus q$ and $p \in \mathcal{U}$, then $q \in \mathcal{U}$. We also require that $\cap \mathfrak{V} \neq \emptyset$, i.e. that agents' assumptions are consistent. Note that there may be some proposition $p \in \mathcal{U}$ such that neither $p \in \mathfrak{V}$ nor $\mathcal{W} \setminus p \in \mathfrak{V}$. That is, an agent may be unaware of a proposition without making any assumption that it is either true or false.

We say that an agent entertains those worlds that the agent does not rule out by assumption:

- (48) Let \mathcal{M} be a background model as defined in (46) and let α be an awareness state as defined in (47). Then $\mathcal{W}_\alpha = \cap \mathfrak{V}$ is set of worlds that an agent ENTERTAINS.

An agent's filtered model is then defined as follows:

- (49) Let \mathcal{M} be a background model as defined in (46) and let α be an awareness state as defined in (47). Then an agent's FILTERED MODEL $\mathcal{M} \upharpoonright \alpha$ is a quadruple $\langle \mathcal{W}', \mathcal{A}', P', \mathcal{U}' \rangle$ where:
- $\mathcal{W}' = \mathcal{W}_\alpha$
 - $\mathcal{A}' = \mathcal{A} \setminus \mathfrak{A}$
 - $P' = P(\cdot | \mathcal{W}_\alpha)$
 - $\mathcal{U}' = \mathcal{U} \upharpoonright (\mathcal{W}' \times \mathcal{A}')$

An agent's filtered model is similar to the agent's background model, but includes only those worlds that agree with the agent's implicit assumptions and only those actions that an agent is aware of. An agent's filtered probability distribution is the agent's background distribution conditioned on those worlds that agree with their implicit assumptions.

Recall that an agent may be unaware of some proposition p without having an implicit assumption about p . Although such an agent may entertain worlds in which p is true and worlds in which p is false, this agent would not be able to distinguish between worlds that *only* differ with respect to the truth value of p . We say that such worlds are equivalent by reason of unawareness and we aggregate worlds that are equivalent by reason of unawareness into states:

- (50) Let α be an awareness state as defined in (47) and let \mathcal{W}_α be a set of worlds that an agent entertains as defined in (48). Two worlds $w, w' \in \mathcal{W}_\alpha$ are EQUIVALENT BY REASON OF UNAWARENESS in α iff for all $p \in \mathcal{P} \setminus \mathcal{U}$, we have $w \in p \leftrightarrow w' \in p$. If w and w' are equivalent by reason of unawareness in α , we write $w \equiv_\alpha w'$.
- (51) S is the set of all equivalence classes on \mathcal{W} induced by \equiv_α . We call each equivalence class $s \in S$ a STATE.

An agent’s decision problem under unawareness and expected utility are defined as follows:

- (52) Let \mathcal{M} be an agent’s background model as defined in (46) and let α be the agent’s awareness state as defined in (47). Then the agent’s DECISION PROBLEM UNDER UNAWARENESS $\delta(\mathcal{M} \upharpoonright \alpha)$ is a quadruple $\langle S, A, \hat{P}, U \rangle$ where:
- a. $S = \mathcal{W}' / \equiv_{\alpha}$
 - b. $A = \mathcal{A}'$
 - c. $\hat{P}(s) = P'(s)$ for all $s \in S$
 - d. $U(s, a) = \sum_{w \in s} P'(w|s) \times \mathcal{U}'(w, a)$ for all $s \in S$ and for all $a \in A$
- (53) The EXPECTED UTILITY of some action a for a decision problem δ is as follows:

$$EU_{\delta}(a) = \sum_{s \in S} \hat{P}(s) \times U(s, a)$$

An agent solves her decision problem by taking the action with the greatest expected utility.

Finally, we must say something about how utterances can affect the awareness states of addressees. In general, we will not be able to give necessary and sufficient conditions for an utterance u to raise an addressee’s awareness of a proposition p due to the highly context-dependent nature of the effects of utterances on awareness states. To illustrate this problem, suppose that Humpty Dumpty has helped Jack move to a new apartment. Bo suggests that Jack send Humpty a fruit basket as a token of his thanks. Jack likes the idea, but Bo knows that he is likely to forget to follow through. In order to remind Jack of the plan, she may utter the following:

- (54) BO: Did you decide on a fruit basket yet?
 JACK: (*Slaps forehead.*)

Bo’s utterance in (54) raises Jack’s awareness of the plan, but her utterance does not explicitly mention anything about sending a fruit basket to Humpty Dumpty. Without knowing the preceding context, we have no hope of understanding what the awareness-related effects of (54) would be. In modeling particular interactions, it is necessary to simply stipulate which propositions and actions an utterance raises awareness of.

Illustrations of agents’ background models and unawareness with and without assumptions are shown in Figures 1, 2, and 3. We can also get a handle on how this model works by reconsidering the context discussed above involving Bo’s search for her keys. We have four worlds, each representing a different location of Bo’s keys: $\mathcal{W} = \{\text{pockets, nightstand, sofa, car}\}$. Likewise, we have one proposition corresponding to each world; \mathcal{P} is the set of the singletons formed from the elements of \mathcal{W} . We also have four actions, each of which corresponds to searching for the keys in a different location: $\mathcal{A} = \{\text{search in pockets, search in nightstand, search in sofa, search in car}\}$. For simplicity, assume Bo has a uniform distribution over the propositions in \mathcal{P} , so $P(\{w\}) = 0.25$ for all $w \in \mathcal{W}$. Finally, let $\mathcal{U}(w, a) = 1$ if there is some location x such that w is a world in which the keys are in x and a is the action of searching for the keys in x . Otherwise, let $\mathcal{U}(w, a) = 0$.

Now consider Bo’s awareness state before Jack speaks up. Since Bo is unaware of the possibility that the keys are in the car, we have $\mathfrak{U} = \{\{\text{car}\}, \mathcal{W} \setminus \{\text{car}\}\}$ and $\mathfrak{A} = \{\text{search in car}\}$. Moreover, she maintains an implicit assumption that the keys are not in the car, so $\mathfrak{V} = \{\mathcal{W} \setminus \{\text{car}\}\}$. This leaves Bo with a filtered model in which $\mathcal{W}' = \mathcal{W}_{\alpha} = \{\text{pockets, nightstand, sofa}\}$, $\mathcal{A}' = \mathcal{A} \setminus \{\text{search in car}\}$, and $P'(\{w\}) = \frac{1}{3}$ for all $w \in \mathcal{W}'$. When we consider Bo’s decision problem, we see that no two worlds in \mathcal{W}' are equivalent for reasons of unawareness, so each state in S will contain exactly one world. \mathcal{U}' is the same as \mathcal{U} , but is only defined for pairs of worlds in \mathcal{W}' and actions in \mathcal{A}' . The expected utility of each action in \mathcal{A}' is $\frac{1}{3}$.

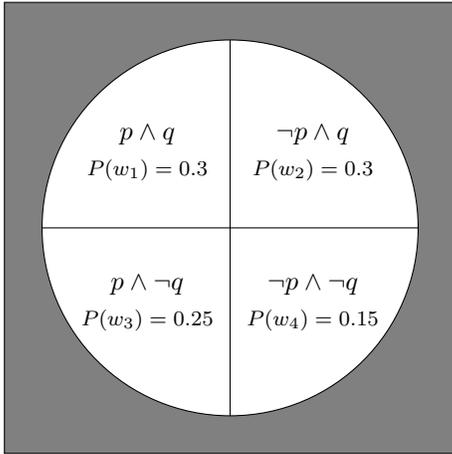


Figure 1: A simple illustration of a background model. There are four worlds with positive probability: w_1, w_2, w_3, w_4 . Each world belongs to one of two two propositions: $p = \{w_1, w_3\}$ and $q = \{w_2, w_4\}$.

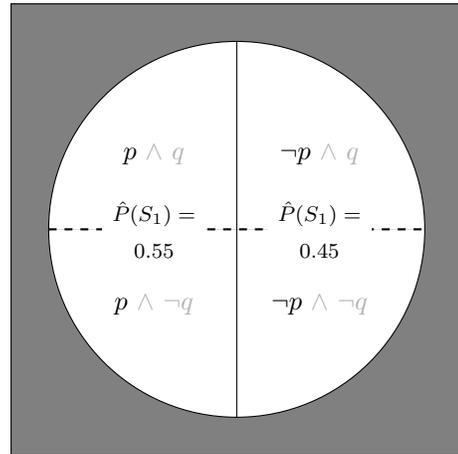


Figure 2: The agent is unaware of the proposition q , but makes no assumption about its truth value. Worlds w_1 and w_3 are aggregated into state S_1 and worlds w_2 and w_4 are aggregated into state S_2 .

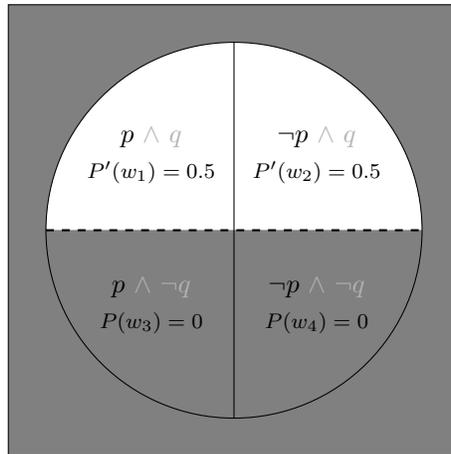


Figure 3: Unawareness with assumptions. The agent is unaware of the proposition q and assumes that q is false. Thus, the agent only entertains two worlds: w_1 and w_2 .

Bo proceeds to search her pockets, her nightstand, and the sofa, since each action has a positive expected utility. When Bo does not find the keys in any of these locations, she is in a conundrum. Given her doxastic state, the keys must be in one of those three locations, but her searches have shown that they are not. Of course, once Jack speaks up, Bo’s awareness state is revised. The propositions $\{\text{car}\}$ and $\mathcal{W} \setminus \{\text{car}\}$ are removed from \mathfrak{U} and \mathfrak{V} , and the action **search in car** is removed from \mathfrak{A} . Having searched her pockets, the nightstand, and the sofa, the only remaining possible worlds is that in which her keys are located in the car. The expected utility of **search in car** will be 1, and Bo will search for her keys in the car.

4.2 The Semantics of Clarity Under Unawareness

Recall that in §3 I gave the following definitions for personal-evidential and simple clarity:

(55) *It is clear to x from e that p is true at a world w and a context c iff for all $i \in x$, i possesses the evidence e at w , i ’s degree of belief in p at w is at least $\mathbf{d}(c)(\llbracket \text{clear} \rrbracket)$, and i ’s possession of e is sufficient grounds for i ’s degree of belief in p to be at least $\mathbf{d}(c)(\llbracket \text{clear} \rrbracket)$.*

(56) *It is clear from e that p is true at a world w and a context c iff*
 $\text{GEN}w' \text{GEN}x (\text{Acc}(w, w') \wedge x \in \text{Rel}(w', c) \rightarrow \text{clear}(x, e, p, w', c))$

The model introduced in the previous section makes it possible to say more precisely what it means for an agent to have a degree of belief greater than some threshold and for a body of evidence to be sufficient grounds for an agent to have a particular degree of belief in a proposition.

On the model introduced above, degrees of belief are represented in terms of a subjective probability distribution. But in this model an agent has two such distributions: one according to the agent’s background model and one according to the agent’s filtered model. Which of these probability distributions should correspond to the degrees of belief that play a role in (55)? To gain some insight into this question, we can consider a situation in which an agent’s background probability distribution and filtered distribution render different judgments about the clarity of some proposition. Suppose Bo and Jack work as job recruiters for a large company and identify résumés to pass on to their boss, Humpty Dumpty. Humpty has never hired a job applicant who majored in linguistics, leading to a mutual belief between Bo and Jack that Humpty has something against linguistics majors and will never hire one. One day, Bo comes across an impressive résumé from a linguistics major, but forgets about Humpty Dumpty’s bias. She shows it to Jack, and the following interaction follows:

(57) BO: This applicant looks good. I think we should pass it up to Humpty.

JACK: It’s not worth our time. It’s clear that Humpty has a bias against linguistics majors.

BO: Oh, right.

Before Jack says anything in (57), it does not seem appropriate to say that Humpty’s bias is clear to Bo. After all, if this were clear to her, she would not have suggested showing Humpty Dumpty the résumé. This suggests that an agent’s filtered probability distribution are relevant for interpreting clarity claims. One might object that on a habitual or generic interpretation, *It is clear to Bo that Humpty Dumpty has a bias against linguistics majors* is true. But note that this is only true if Bo does not regularly operate with an implicit assumption that Humpty does not have such a bias. That is, even on the interpretation according to which it *is* clear to Bo that Humpty Dumpty has a bias against linguistics majors, this interpretation still relies on considering what Bo habitually believes according to her *filtered* probability distribution. Thus, it seems that an agent’s filtered distribution, which represents an agent’s beliefs under unawareness, is crucial for determining what is clear to that agent.

Next, what does it mean for a body of evidence e to be sufficient grounds for an agent x 's degree of belief in some proposition p to exceed some threshold θ ? As a first pass, we could assume something along the following lines:

- (58) Let x be an agent with a filtered subjective probability distribution P' . Let E be a set of propositions representing x 's total body of evidence. We say that some body of evidence $e \subseteq E$ is sufficient grounds for x 's subjective probability in p to be greater than θ iff $P'(p|E) > \theta$ and $P'(p|E \setminus e) \not> \theta$.

That is, a body of evidence e is sufficient grounds for an agent's degree of belief in p to exceed some threshold θ if the agent's degree of belief in p exceeds θ given the agent's total body of evidence (including e), but would not exceed θ if e were removed from the agent's total body of evidence. Note that this does not preclude there being some other form of evidence, e' , the removal of which from the agent's total body of evidence would cause the agent to assign to p an even lower degree of belief than the removal of e would cause the agent to assign to p . In addition, this definition allows for there to be some additional evidence e'' that is not currently part of the agent's total body of evidence, but which would further increase the agent's degree of belief in p were it to be included in the agent's total body of evidence.

Incorporating this into our existing semantics for personal-evidential clarity gives the following:

- (59) Truth conditions for personal-evidential clarity (second version, to be revised):
 Let i be an agent with filtered subjective probability distribution in world w and context c $P_i^{w,c}$. Let $E_i^{w,c}$ be i 's total body of evidence in w and c , and let $\theta_{clear}^{w,c}$ be the contextually relevant threshold for clarity. Then *It is clear to x from e that p is true in w and c* iff:
 $\forall i \in x (P_i^{w,c}(p|E_i^{w,c}) > \theta_{clear}^{w,c} \wedge P_i^{w,c}(p|E_i^{w,c} \setminus e) \not> \theta_{clear}^{w,c})$

The definition for evidential clarity remains the same, with the definition in (59) used to interpret $clear(x, e, p, w, c)$.

Unfortunately, (59) faces its own problems.²⁸ First, note that it does not seem particularly odd for a speaker to make multiple assertions of personal clarity with the same prejacent, but different bodies of evidence:

- (60) BO: Why do you think it's clear that Dostoyevsky is a genius?
 JACK: It's clear to me from reading *The Brothers Karamazov*. And it's also clear to me from reading *Crime and Punishment*.

The proposal in (59) predicts Jack's statement in (60) to be true only if reading both *The Brothers Karamazov* and *Crime and Punishment* were necessary for Jack to come to believe that Dostoyevsky is a genius. This is because for Jack's statement to be true, it must be the case that removing evidence related to *The Brothers Karamazov* from Jack's total body of evidence, but leaving evidence related to *Crime and Punishment* in his total body of evidence, would lead him to no longer believe that Dostoyevsky is a genius. The same must also be true if we were to remove evidence related to *Crime and Punishment*, but keep evidence related to *The Brothers Karamazov*. Thus, both pieces of evidence would be needed for Jack to believe that Dostoyevsky is a genius. But Jack's utterance also seems compatible with an interpretation such that he would believe Dostoyevsky to be a genius had he only read *either* of the two novels.

To illustrate another problem with (59), suppose we are in a world $w_{Dostoyevsky}$ in which all relevant individuals have read the complete works of Dostoyevsky. Further suppose that reading

²⁸My thanks to Deniz Rudin (p.c.) for drawing my attention to this issue.

any work by Dostoyevsky is sufficient evidence for concluding that Dostoyevsky is a genius. In such a world, it seems plausible that a speaker could assert the following felicitously:

(61) It's clear from reading *Crime and Punishment* that Dostoyevsky is a genius.

However, the proposal in (59) predicts (61) to be false in the context described. First, note that according to (59), it is not clear to anyone in $w_{Dostoyevsky}$ from reading *Crime and Punishment* that Dostoyevsky is a genius. This is because the individuals in $w_{Dostoyevsky}$ would believe Dostoyevsky to be a genius even if they hadn't read *Crime and Punishment*. Next, since there is no explicit experiencer argument in (61), its truth value depends upon the beliefs of relevant individuals in general in worlds accessible from $w_{Dostoyevsky}$. A defender of (59) might point out that in some of these worlds, it may be the case that the relevant individuals have not all read the complete works of Dostoyevsky. In those worlds, it would be true according to (59) that it is clear to the relevant individuals from reading *Crime and Punishment* that Dostoyevsky is a genius. Unfortunately, to save (59) it is not enough for there to be *some* worlds accessible from $w_{Dostoyevsky}$ where reading *Crime and Punishment* makes it clear to the relevant people that Dostoyevsky is a genius. Rather, it must be the case that this is true about worlds accessible from $w_{Dostoyevsky}$ *in general*, which seems implausible.

To resolve these issues, I propose the following adjustment. Rather than saying that the evidence e must cause the experiencer to believe p relative to the rest of the experiencer's evidence, $E \setminus e$, we say that e must cause the experiencer to believe p relative to some core subset of E :

(62) Truth conditions for personal-evidential clarity (final version):

Let i be an agent with filtered subjective probability distribution in world w and context c $P_i^{w,c}$. Let $E_i^{w,c}$ be i 's total body of evidence in w and c , and let $\mathcal{E}_i^{w,c} \subseteq E_i^{w,c}$ be some core set of i 's evidence in w and c . Let $\theta_{clear}^{w,c}$ be the contextually relevant threshold for clarity. Then *It is clear to x from e that p is true in w and c iff:*
 $\forall i \in x (P_i^{w,c}(p|\mathcal{E}_i^{w,c}) > \theta_{clear}^{w,c} \wedge P_i^{w,c}(p|E_i^{w,c} \setminus e) \not> \theta_{clear}^{w,c})$

Note that the definition in (62) does not tell us what this core evidence consists in. One way to think of core evidence is as the evidence that we can assume a "normal" reasoner would possess.²⁹ In general, what will count as core evidence will be highly context dependent. In the examples focusing on Dostoyevsky, this evidence will include propositions relating to the human condition necessary for appreciating great works of literature but will likely exclude evidence about Dostoyevsky's work in particular. In example (40), this evidence will include the propositions regarding the rules of chess.

4.3 Clarity Assertions & Awareness

Let's now consider how this perspective on clarity assertions provides an explanation for the function of such assertions, even when they appear to be uninformative. We return to our original example in which Jack is looking for pastries and Bo is accompanying him. When they arrive outside Tartine and Arizmendi, Tartine has a long line outside, but Arizmendi does not. Bo then utters the following:

(63) BO: It's clear that there is a wait at Tartine.

²⁹This way of thinking about how some body of evidence causes an agent to have a particular belief is not dissimilar from philosophical approaches to causation according to which an event causes some result if that result would not obtain in the most normal counterfactual situations in which the causal event did not take place (Lewis 1973; Halpern and Pearl 2005; Hall 2007; Hitchcock 2007; Hitchcock and Knobe 2009).

The awareness model introduced above first allows us to resurrect Barker and Taranto’s (2003) claim that an assertion of the clarity of p converts mutual belief in p to common belief in p . Suppose that before Bo utters (63), Bo and Jack have seen the line outside Tartine and conclude from this evidence that it is extremely likely that there is a wait at Tartine. That is, it is mutually believed that there is a wait at Tartine. Due to unawareness, Jack may have failed to consider the issue of whether Bo believes that there is a wait at Tartine, perhaps even maintaining an implicit assumption that Bo does not believe there is a wait at Tartine. In this case, it would not be commonly believed that there is a wait at Tartine. By uttering (63), Bo makes Jack aware that she believes there is a wait at Tartine, thereby adding the proposition that there is a wait at Tartine to the common ground.³⁰

The awareness model also allows us to consider the possibility that there being a wait at Tartine is not even a mutual belief. Due to unawareness, a discourse participant may fail to explicitly believe p despite having access to the evidence necessary to conclude p and having the inferential resources required to conclude p on the basis of the evidence. A failure on the part of an agent to explicitly believe p in a context that would, under full awareness, cause the agent to believe p may take several different forms. An agent may fail to explicitly believe p while maintaining an implicit assumption that p is false, but an agent may also fail to explicitly believe p while making no such assumption. In addition, an agent who fails to explicitly believe p may be unaware of the evidence supporting p or may be aware of this evidence, but unaware of what follows from it. Below, I consider each of these possibilities.

To model these possibilities using Bo and Jack’s interaction as an example, first let $\mathcal{W} = \{w_1, w_2, w_3, w_4, w_5, w_6, w_7, w_8\}$. Jack’s decision problem will depend on a number of factors. For the sake of simplicity, I limit the issues considered here to whether there is a wait at either Arizmendi or Tartine and which bakery makes better pastries. Let $\mathcal{P} = \{p, q, r\}$, where $p = \{w_1, w_3, w_5, w_7\}$ is the proposition that there is a wait at Tartine, $q = \{w_1, w_2, w_3, w_4\}$ is the proposition that there is a wait at Arizmendi, and $r = \{w_1, w_2, w_5, w_6\}$ is the proposition that Tartine has better pastries than Arizmendi. Let $\mathcal{M} = \langle \mathcal{W}, \mathcal{A}, P, \mathcal{U} \rangle$ be Jack’s background model. Let $\mathcal{A} = \{\text{Tartine}, \text{Arizmendi}\}$, where **Tartine** is the action of going to Tartine and **Arizmendi** is the action of going to Arizmendi. In defining Jack’s utility function, I take it that Jack would prefer to not wait for his pastries and would prefer to get the best pastries. However, if a single action cannot accomplish both goals, he prefers not waiting over getting better pastries. I define \mathcal{U} as follows:

$$U(a, w) = \begin{cases} 1 & \text{if } (a = \text{Tartine} \wedge w \in (\mathcal{W} \setminus p) \cap r) \vee (a = \text{Arizmendi} \wedge w \in (\mathcal{W} \setminus q) \cap (\mathcal{W} \setminus r)) \\ 0.75 & \text{if } (a = \text{Tartine} \wedge w \in (\mathcal{W} \setminus p) \cap (\mathcal{W} \setminus r)) \vee (a = \text{Arizmendi} \wedge w \in (\mathcal{W} \setminus q) \cap r) \\ 0.5 & \text{if } (a = \text{Tartine} \wedge w \in p \cap r) \vee (a = \text{Arizmendi} \wedge w \in q \cap (\mathcal{W} \setminus r)) \\ 0.25 & \text{if } (a = \text{Tartine} \wedge w \in p \cap (\mathcal{W} \setminus r)) \vee (a = \text{Arizmendi} \wedge w \in q \cap r) \end{cases}$$

The definitions for clarity in (62) make reference to conditioning on some contextually relevant, core body of evidence. So rather than define P directly, I consider $P(\cdot|\mathcal{E})$, where \mathcal{E} is Jack’s core body of evidence in the context. For the moment, let’s assume that \mathcal{E} contains the evidence of the line outside Tartine. Since \mathcal{E} includes the proposition that there is a line outside Tartine, $P(p|\mathcal{E})$, i.e. the probability that there is a wait at Tartine given Jack’s total body of evidence, should be relatively large. We also assume that \mathcal{E} contains the proposition that there is no line outside Arizmendi, so $P(q|\mathcal{E})$ should be relatively small. Finally, let’s assume that Jack thinks that Tartine probably has better pastries than Arizmendi, so $P(r|\mathcal{E})$ is large.

$$P(\varphi|\mathcal{E}) = \begin{cases} 0.75 & \text{if } \varphi = p, r \\ 0.25 & \text{if } \varphi = q \end{cases}$$

³⁰As we’ve seen, it is possible for a speaker to assert the clarity of some proposition p without believing in p , although this requires the speaker to be committed to having abnormal beliefs. If we assume that Bo takes herself to have normal beliefs, then her assertion in (63) does commit her to believing that there is a wait at Tartine.

We next consider Jack’s awareness state $\alpha = \langle \mathfrak{U}, \mathfrak{A}, \mathfrak{B} \rangle$. We will suppose at first that Jack is unaware of the possibility of a wait at either Tartine or Arizmendi, so $\mathfrak{U} = \{p, q\}$. We will also suppose for this first example that Jack makes the implicit assumptions that there is no wait at either Tartine or Arizmendi. Thus, $\mathfrak{B} = \{\mathcal{W} \setminus p, \mathcal{W} \setminus q\}$. We’ll take Jack to be fully aware of the actions of going to Tartine and going to Arizmendi, so $\mathfrak{A} = \emptyset$. In Jack’s filtered model $\mathcal{M} \upharpoonright \alpha$, we have $\mathcal{W}_\alpha = \{w_6, w_8\}$ and $\mathcal{A}' = \mathcal{A}$. We still have $P'(r|\mathcal{E}) = 0.75$, although $P'(p|\mathcal{E}) = P'(q|\mathcal{E}) = 0$ given Jack’s implicit assumption that there is no wait at either Arizmendi or Tartine. Jack’s filtered utility function \mathcal{U}' is defined as follows:

$$\mathcal{U}'(a, w) = \begin{cases} 1 & \text{if } (a = \text{Tartine} \wedge w \in p) \vee (a = \text{Arizmendi} \wedge w \in \mathcal{W} \setminus p) \\ 0.75 & \text{if } (a = \text{Tartine} \wedge w \in \mathcal{W} \setminus p) \vee (a = \text{Arizmendi} \wedge w \in p) \end{cases}$$

Let δ be Jack’s decision problem. Since Jack is only aware of two worlds, w_6 and w_8 , and since each is distinguished by a proposition that Jack is aware of, namely r , then there are two states relevant to Jack’s decision problem: one corresponding to w_6 and one corresponding to w_8 . With P' and \mathcal{U}' , we calculate the expected utility of each action as follows:

$$EU_\delta(\text{Tartine}) = 0.75 \times 1 + 0.25 \times 0.75 = 0.9375$$

$$EU_\delta(\text{Arizmendi}) = 0.75 \times 0.25 + 0.25 \times 1 = 0.4375$$

Thus, we expect Jack to choose to go to Tartine despite the fact that it has a long line.

Now consider the effect of Bo’s assertion that it is clear there is a wait at Tartine. On the proposal offered here, her assertion is true so long as, in general, relevant individuals would assign a sufficiently high subjective probability to p . Even if, in the actual world, Jack is making an implicit assumption that p is false, relevant individuals would, in general, not make such an assumption. Thus, p is clear in the context, despite the fact that Jack is unaware of p and his behavior reflects an assumption that p is false. Bo’s utterance changes the situation by immediately raising Jack’s awareness of p . We can also assume that this utterance raises his general awareness of propositions relating to waits at bakeries, so that he becomes aware of q as well. Since there are no propositions of which Jack is unaware, his filtered model is now the same as his background model. Let δ' be Jack’s decision problem under his new awareness state. Calculating expected utility of each action gives the following values:

$$EU_{\delta'}(\text{Tartine}) \approx 0.621$$

$$EU_{\delta'}(\text{Arizmendi}) \approx 0.689$$

As expected, after Bo’s assertion, going to Arizmendi is the preferred action.

The above example rests on a number of assumptions that do not reflect all cases in which clarity is asserted. We took Jack’s core body of evidence \mathcal{E} to include the fact that there was a line outside Tartine. We also took Jack to hold implicit assumptions that p and q were false. We now consider circumstances in which one or both of these assumptions do not hold. First, we will consider the same situation described above but now suppose that \mathcal{E} does not contain the information that there is a line outside Tartine and no line outside Arizmendi. For whatever reason, Jack fails to recognize the presence of a line at Tartine and an absence of a line at Arizmendi.³¹

³¹This may seem to be an unreasonable assumption, since it is hard to imagine how Jack could fail to notice the line outside Tartine and absence of line at Arizmendi. But note that, in general, it is not unreasonable to think that agents could fail to recognize evidence available to them. Suppose Bo and Jack are looking at a photograph of a woman and are trying to learn as much as they can about her from the photograph. The woman in the photograph happens to be writing with her left hand, and Bo utters, “It’s clear that she’s left-handed.” It is plausible that Jack could have failed to notice the relevant evidence supporting Bo’s assertion. He simply may not have been focusing on which hand the woman was writing with as he scrutinized the photograph.

Suppose that Jack’s subjective probability distribution is now the following:

$$P(\varphi|\mathcal{E}) = \begin{cases} 0.75 & \text{if } \varphi = r \\ 0.50 & \text{if } \varphi = p, q \end{cases}$$

We will again assume that Jack makes the implicit assumption that both p and q are false. Note that because Jack has failed to recognize the evidence, he no longer counts as a relevant experiencer in the actual world, and so his beliefs are irrelevant to the truth of Bo’s assertion. When Bo utters (63), Jack’s assumptions are overturned, but this is not the end of the story. If these assumptions are overturned and there is no other change to Jack’s mental state, he will continue to assign the proposition that there is a wait at Tartine probability 0.5. Of course, if Jack accepts Bo’s assertion, it is unlikely that his own beliefs about p and q will remain unchanged. In particular, upon accepting Bo’s assertion, the subjective probability that Jack assigns to p should be revised upward, after which going to Arizmendi will be the preferred action.³²

Note that in accepting Bo’s assertion, Jack is also likely to become aware that there is a line outside Tartine. Bo’s assertion entails that there is some body of evidence that causes some group of experiencers to assign a high probability to the proposition that there is a wait at Tartine. Jack will naturally interpret Bo as being one, and perhaps the only, experiencer for whom this is true in the actual world. As we saw in the discussion of informative clarity assertions, it is possible that a speaker may assert simple clarity when the evidence is not available to the addressee.³³ However, in most circumstances, Bo will not possess private evidence for her belief that there is a wait at Tartine, and Jack will consider whether her assertion may be based on some contextually available evidence. This triggers Jack to reconsider the contextually available evidence, in which case he should notice the line outside Tartine. Jack may also notice the absence of a line outside Arizmendi, which in turn will cause the subjective probability he assigns to q to be revised downward.

So far, we have only considered cases in which Jack assumes that the proposition whose clarity is asserted is false. However, our awareness model allows us to consider a case in which an addressee makes no implicit assumptions about the proposition’s truth value.³⁴ Suppose that Jack is unaware of the propositions that there is a wait at Tartine or Arizmendi, but does not make any implicit assumptions that these propositions are false. Assume for now that Jack’s core evidence \mathcal{E} contains the information that there is a line outside Tartine. Then Jack’s filtered model will be as described above, with the exception that for his awareness state we have $\mathfrak{U} = \{p, q\}$ and $\mathfrak{V} = \emptyset$.

Since there are no worlds that Jack rules out by reason of unawareness, his filtered distribution P' is the same as his distribution under full awareness P . In particular, this means that $P'(p|\mathcal{E}) = 0.75$ and $P'(q|\mathcal{E}) = 0.25$, i.e. that he still believes it is likely that there is a wait at Tartine, but unlikely that there is a wait at Arizmendi. This ends up being reflected in Jack’s decision problem. Since Jack is unaware of p and q , the worlds under consideration are aggregated into two states: one in which r is true and one in which r is false. Let $S = \{s_1, s_2\}$, where s_1 is the state in which r is true and s_2 is the state in which r is false. We have $\hat{P}(s_1|\mathcal{E}) = 0.75$ and

³²Again, it is technically possible for an agent to accept the clarity of p without believing p , but this is unexpected in normal circumstances.

³³In Bo and Jack’s case, we could imagine a scenario in which there is no visible line outside Tartine, but Bo walks in to see what it looks like inside. Upon entering, she sees that the building is packed. She comes back outside and announces to Jack that it is clear there is a wait at Tartine. In this case, her assertion is felicitous despite being based on privately held evidence.

³⁴Yet another option is that the addressee may assume that the prejacent is true. But this case will usually not be of interest. If an agent makes an implicit assumption that some proposition p is true, the agent’s behavior will be consistent with a belief in p . If this implicit assumption is overturned, the agent will continue to act as if p were very likely. Thus, asserting the clarity of p to someone who assumes that p is true will likely have few consequences.

$\hat{P}(s_2|\mathcal{E}) = 0.25$. We have the following utility function on state, action pairs:

$$U(s, w) \approx \begin{cases} 0.470 & \text{if } s = s_1 \\ 0.212 & \text{if } s = s_2 \wedge a = \text{Arizmendi} \\ 0.141 & \text{if } s = s_2 \wedge a = \text{Tartine} \end{cases}$$

That is, if we are in s_1 , the utility of going to Tartine and going to Arizmendi are the same. If we are in s_2 , then going to Arizmendi has a higher utility. Thus, Jack will choose to go to Arizmendi even though he is unaware that there is a wait at Tartine.

This result may seem unintuitive, based on the following line of reasoning. If Jack is unaware of the proposition that there is a wait at Tartine, then he should go to Tartine; after all, they have the better pastries. But so long as Jack possesses the evidence regarding lines outside Tartine and Arizmendi and does not make any implicit assumptions about the waits at the two bakeries, he will implicitly believe that there is a wait at Tartine. Bo's assertion that it's clear there is a wait at Tartine cannot have any effect on Jack's decision problem. Although unintuitive at first glance, this result should not be surprising given how we have modelled the situation. As described, Jack recognizes that there is a line outside Tartine, and might respond to the situation accordingly:

(64) JACK: Oh, there's a huge line outside Tartine. I had better go to Arizmendi.

Jack might not be aware of the proposition that there is a wait at Tartine, and therefore would not assert this. But he does believe that there is a line outside Tartine and understands the consequences of this for his decision problem. In such a situation, Bo's utterance does seem like it would be superfluous.

Finally, we consider the case is one in which Jack is unaware of the proposition that there is a wait at Tartine, makes no implicit assumption about this proposition, but also has failed to incorporate the evidence about the line into his core evidence \mathcal{E} . Such a situation would be highly similar to the situation described above in which Jack fails to recognize the evidence and makes an implicit assumption that p is false. Bo's utterance would raise Jack's awareness of the proposition that there is a wait at Tartine. In accepting Bo's assertion, Jack would likely come to believe that there is a wait at Tartine. In addition, without any reason to suspect that Bo had some private knowledge about waits at Tartine, Jack would conclude that she based her assertion on contextually available evidence. He would then reconsider the evidence available to him, updating his evidence with the fact that there is a line outside Tartine.

Before moving on, it is important to revisit several examples discussed above which proved problematic for previous analyses of clarity statements. First, recall that Bronnikov (2008) was able to account for repeated assertions of clarity based on similar bodies of evidence. For example, suppose a woman walks by wearing a white coat and a stethoscope, after which Bo utters (65a). Immediately afterwards a man walks by dressed in the same way and Jack utters (65b).

- (65) a. BO: It's clear that she's a doctor.
 b. JACK: It's clear that he's a doctor too.

Let p denote the proposition that the woman is a doctor and let q denote the proposition that the man is a doctor. Let e denote the proposition that the woman is wearing a white coat and stethoscope and let f denote the proposition that the man is wearing a white coat and stethoscope. Nothing precludes either Bo or Jack from being unaware of p or q or of having failed to accept the evidence e or f . Thus, if Bo supposes that Jack is unaware of p or e , she may utter (65a) to raise his awareness of either of these propositions. Likewise, Jack may utter (65b) to raise Bo's awareness of q or f .

Bronnikov also points out that clarity assertions are generally infelicitous when they involve “trivial” inferences (66).

- (66) a. Mindy is a doctor.
b. # Therefore, it is clear that Mindy is a doctor.

On the account offered here, (66b) is, strictly speaking, true. Any relevant experiencer will believe that Mindy is a doctor on the basis of the evidence that Mindy is a doctor. The oddness of asserting (66b) according to the present account is that (66a) has already raised awareness of the proposition that Mindy is a doctor before (66b) is uttered. For this reason, (66b) serves no purpose; it is neither informative, nor awareness-raising. Thus, we expect it to be infelicitous.

Bronnikov also raises the issue of asserting the clarity of necessarily true propositions (67).

- (67) Take a number n that is divisible by 9. Then it is clear that n is divisible by 3.

Here, we can see the role of the clarity assertion as that of raising awareness of the proposition that n is divisible by 3, which discourse participants may have been unaware of despite knowing n to be divisible by 9. One could ask what purpose raising awareness would serve in this case, since any agent who believes that n is divisible by 9 must, at least implicitly, believe that n is divisible by 3. Stated differently, we cannot imagine discourse participants to hold an implicit assumption that n is not divisible by 3 once they know that n is divisible by 9. In this case, what is the purpose of raising awareness that n is divisible by 3?

A similar issue arises if we consider (26), repeated below as (68), which was argued to be problematic for Bronnikov’s proposal.

- (68) It is clear that we have a highly polarized, very sharp, and I think at times, too sharp level of anger expressed across the political aisles in this country.

It was noted above that the purpose of asserting clarity in (68) appears to be that of raising awareness of a proposition that the addressee implicitly believed, but was unaware of. The present account can make sense of such a change in awareness, but (68) is a case in which the addressee is not obviously making an implicit assumption that the prejacent of the clarity assertion is false. Again, we face the question of why it is worth raising awareness of a proposition even if the addressee does not make an implicit assumption that the proposition is false.

To explain such cases, I note that it may be important to raise awareness of particular propositions for discourse structural reasons. For example, (68) could be used to raise awareness of a proposition already implicitly believed by all discourse participants in order to establish it as a discourse topic worthy of further elaboration. The clarity statement in (67) might serve to conclude a discourse, e.g. as the end of a long proof whose purpose is to show that a number n with certain properties is divisible by 3.³⁵

One way to capture these intuitions is to extend the decision theoretic model by representing inferences as a type of action. Discourse participants may take as a goal the performance of particular inferences.

- (69) Let \mathcal{A} be the set of actions in an agent’s background model as defined in (46). Then $\mathcal{I} \subseteq \mathcal{A}$ is the set of the agent’s INFERENCE ACTIONS.

³⁵Early work on computational representations of discourse structure often emphasized the importance of attention and awareness for coherently structuring discourse (Grosz and Sidner 1986; Walker 1996), even though more recent work on discourse structure (Asher and Lascarides 2003; Kehler 2005, 2011) has placed less explicitly emphasis on the importance of awareness.

We can think of an inferential action as an action related to a set of premises and a conclusion, all of which are propositions.

- (70) Let $i \in \mathcal{I}$ be an agent’s inferential action as defined in (69) and let \mathcal{W} be the set of worlds in the agent’s background model as defined in (46). Then $\Pi(i) \subseteq \wp(\mathcal{W})$ is the set of PREMISES of i and $K(i) \in \wp(\mathcal{W})$ is the CONCLUSION of i .

Next, we leverage the fact that agents may be unaware of actions as well as propositions. In the case of inferential actions, it seems plausible to assume that an agent is aware of the inferential action only if the agent is aware of both its premise and conclusion:

- (71) Let $i \in \mathcal{I}$ be an inferential action in some agent’s background model as defined in (69). Let $\alpha = \langle \mathfrak{U}, \mathfrak{V}, \mathfrak{A} \rangle$ be the agent’s awareness state as defined in (47). If there is some $p \in \Pi(i) \cup K(i)$ such that $p \in \mathfrak{U}$, then $i \in \mathfrak{A}$.

On this view, the role of (67) may be that of making the addressee aware of some proposition that serves as the conclusion to an inference, thus making that inference possible. In contrast, (68) may serve to make the addressee aware of a premise necessary for performing a future inference.

The idea of treating inferences as actions obviously parallels Bronnikov’s proposal about clarity, according to which inference is a type of action. Recent work in philosophy and psychology has also treated inference as a type of action (Icard 2014; Vul et al. 2014). However, the similarities to Bronnikov (2008) only go so far. For Bronnikov, inferential actions are involved in the truth conditions of *clear*. In the current proposal, clarity assertions can raise awareness of propositions that agents must be aware of in order to perform particular inferences. Crucially, the present proposal also allows clarity assertions to raise awareness of propositions that are relevant for other, non-inferential reasons.

5 Conclusion

I have offered a new proposal for the semantics of clarity statements, the most notable feature of which is its approach to the interpretation of implicit experiencer arguments of *clear*. On this proposal, clarity statements with implicit experiencer arguments are interpreted as generically quantifying over some set of contextually relevant individuals. This proposal allows for clarity statements to be both true and informative, particularly when the set of contextually relevant individuals does not include discourse participants. Still, there remain cases in which clarity assertions appear to be uninformative. In these cases, clarity assertions may play the role of raising awareness of propositions that are supported by evidence available to discourse participants, but which these participants may fail to recognize due to inattention. Raising awareness can, in turn, have downstream effects including, but not limited to, effects on addressees’ resolutions to decision problems.

The second of these two resolutions to the paradox of asserting clarity relies on a very general understanding of how assertions of clarity can affect an individual’s awareness state. This raises the following question: why would a speaker assert the clarity of a proposition in order to raise awareness of a proposition, rather than something else? Any mention of a proposition by a speaker can raise an addressee’s awareness of that proposition, so why does the speaker not mention the proposition in some other way? In fact, other ways of mentioning the relevant proposition *can* raise awareness. Consider Bo and Jack on their pastry expedition. In the original example, Bo asserted that it was clear that there would be a wait at Tartine, but any of the following would also have the effect of raising Jack’s awareness of the wait at Tartine:

- (72) a. BO: Well, there’s going to be a wait at Tartine.

- b. BO: It looks like there's going to be a wait at Tartine.
- c. BO: How long do you think we'll have to wait in that line at Tartine?
- d. BO: Needless to say, there is a wait at Tartine.
- e. BO: It's obvious that there is a wait at Tartine.

One may object at this point that the possibility of raising awareness through expressions other than clarity statements leaves the puzzle of asserting clarity unresolved. Since the awareness-raising potential of clarity assertions has nothing to do with *clear* in particular, there must be more to explain. But this objection is misguided in that my aim was to show that there is a function served by assertions of clarity when they appear to be uninformative, not that there is a function that is *uniquely* served by assertions of clarity.

Nonetheless, it is possible to identify functions served by asserting clarity that are not served by other awareness-raising expressions. One of the key features of the semantic analysis of *clear* presented here is that if p is clear, then there exists some body of evidence that supports concluding p . As discussed in the previous section, when the addressee concludes that this evidence is not privately held by the speaker, this can motivate the addressee to reconsider the contextually available evidence. Thus, clarity assertions well-suited to draw addressees' attention to evidence in the context supporting the prejacent of *clear*, leading the addressees to draw the relevant conclusion on her own.

Awareness-raising assertions of the clarity of p can also serve the following function that would not be served by an outright assertion of p . On the proposal put forward here, if a proposition is clear simpliciter, it is something that relevant individuals would (in general) conclude. If we suppose that discourse participants are included among these relevant individuals, then a clarity statement entails that these discourse participants could be expected to conclude that the prejacent is true if they had access to the relevant information or if they considered the issue explicitly. Taking this reasoning one step further, suppose that an assertion of the clarity of p is in pragmatic competition with an outright assertion of p , such as (72a). Assuming the addressee is a relevant experiencer, the clarity assertion is more informative than a bare assertion of p in that the clarity assertion communicates the speaker's belief that the addressee is capable of concluding p . In contrast, the outright assertion of p does not communicate this. If a speaker asserted p , rather than the clarity of p , an addressee might reason that the speaker opted for the less informative utterance (p), rather than the more informative utterance (*clear*(p)), because the speaker believed that the more informative utterance was false in virtue of the addressee being incapable of concluding p on her own. In other words, when an assertion of the clarity of p is in pragmatic competition with an assertion of p , the latter may implicate that the speaker does not take the addressee to be capable of concluding p . Of course, such an implicature could be quite face-threatening towards the addressee. In order to avoid this face-threatening implicature, a clarity assertion may be used instead of an outright assertion of p in order to raise an addressee's awareness of p while still implying that the addressee is capable of concluding p .³⁶

In closing, I note that seemingly uninformative clarity assertions reveal deficiencies in treating linguistic agents as perfectly rational and logically omniscient. There are undoubtedly many more linguistic phenomena that can demonstrate the deficiencies in such assumptions, and we will be able to make better sense of these phenomena by taking seriously psychological realities of discourse participants, including constraints on their attentional and memory capacities.

³⁶Similar reasoning can be extended to other expressions that entail or imply that the addressee already believes p , e.g. *We all know that p, As you already know, p*. Thus, clarity assertions' ability to avoid this particular face-threatening implicature is not unique to clarity assertions. Nonetheless, this aspect of clarity assertions does distinguish them from bare assertions of the prejacent.

References

- Anand, P. (2008). Predicates of taste: Context-shift, assessment-shift, or something else? Presented at the 2nd International Conference on Quotation and Meaning, ZAS, Berlin.
- Asher, N. and Lascarides, A. (2003). *Logics of Conversation*. Cambridge University Press, Cambridge.
- Barker, C. (2009). Clarity and the grammar of skepticism. *Mind & Language*, 24(3):253–273.
- Barker, C. (2011). Commentary on Wolf and Cohen: Reasoning about public evidence. In Égré, P. and Klinedinst, N., editors, *Vagueness and Language Use*, Palgrave Studies in Pragmatics, Language and Cognition Series, pages 191–203. Palgrave Macmillan, New York.
- Barker, C. and Taranto, G. (2003). The paradox of asserting clarity. In Koskinen, P., editor, *Proceedings of the Western Conference on Linguistics (WECOL) 2002*, volume 14, pages 10–21, Fresno, CA. Department of Linguistics, California State University.
- Bledin, J. and Rawlins, K. (2016). Epistemic resistance moves. In Moroney, M., Little, C.-R., Collard, J., and Burgdorf, D., editors, *Proceedings of SALT 26*, pages 620–640.
- Bronnikov, G. (2008). The paradox of clarity: Defending the missing inference theory. In Friedman, T. and Ito, S., editors, *Proceedings of SALT 18*, pages 144–157, Ithaca, NY. Cornell University.
- Ciardelli, I., Groenendijk, J., and Roelofsen, F. (2011). Attention! *Might* in Inquisitive Semantics. In Cormany, E., Ito, S., and Lutz, D., editors, *Proceedings of SALT 19*, pages 91–108.
- Condoravdi, C. and Gawron, J. M. (1996). The context-dependency of implicit arguments. In Kanazawa, M., Piñon, C., and de Swart, H., editors, *Quantifiers, Deduction, and Context*, chapter 1. CSLI Publications, Stanford.
- Davies, M. (2008). The corpus of contemporary American English: 450 million words, 1990-present. Available online at <http://corpus.byu.edu/coca/>.
- Davis, C. (2009). Decisions, dynamics, and the japanese particle *yo*. *Journal of Semantics*, 26(4):329–366.
- de Jager, S. T. (2009). *Now that you mention it, I wonder . . .*. PhD thesis, Universiteit van Amsterdam, Amsterdam.
- Duc, H. N. (2001). *Resource-Bound Reasoning about Knowledge*. PhD thesis, University of Leipzig.
- Fagin, R. and Halpern, J. Y. (1987). Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34(1):39–76.
- Farkas, D. F. and Bruce, K. B. (2010). On reacting to assertions and polar questions. *Journal of Semantics*, 27(1):81–118.
- von Fintel, K. and Gillies, A. S. (2010). Must . . . stay . . . strong! *Natural Language Semantics*, 18(4):351–383.
- Franke, M. and de Jager, T. (2007). The relevance of awareness. In Aloni, M., Dekker, P., and Roelofsen, F., editors, *Proceedings of the Sixteenth Amsterdam Colloquium*, pages 91–96.

- Franke, M. and de Jager, T. (2011). Now that you mention it: Awareness dynamics in discourse and decisions. In Benz, A., Ebert, C., Jäger, G., and van Rooij, R., editors, *Language, Games, and Evolution*, pages 60–91. Springer-Verlag, Berlin.
- Ginzburg, J. (2012). *The Interactive Stance: Meaning for Conversation*. Oxford University Press, Oxford.
- Grosz, B. J. and Sidner, C. L. (1986). Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204.
- Gunlogson, C. (2001). *True to Form: Rising and Falling Declaratives as Questions in English*. PhD thesis, University of California, Santa Cruz.
- Gunlogson, C. (2008). A question of commitment. *Belgian Journal of Linguistics*, 22:101–136.
- Hall, N. (2007). Structural equations and causation. *Philosophical Studies*, 132(1):109–136.
- Halpern, J. Y. and Pearl, J. (2005). Causes and explanations: A structural-model approach. part i: Causes. *The British journal for the philosophy of science*, 56(4):843–887.
- Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. PhD thesis, University of Massachusetts, Amherst.
- Hitchcock, C. (2007). Prevention, preemption, and the principle of sufficient reason. *The Philosophical Review*, 116(4):495–532.
- Hitchcock, C. and Knobe, J. (2009). Cause and norm. *The Journal of Philosophy*, 106(11):587–612.
- Icard, T. (2014). *The Algorithmic Mind: A Study of Inference in Action*. PhD thesis, Stanford University, Stanford.
- Kehler, A. (2005). Discourse coherence. In Horn, L. R. and Ward, G., editors, *The Handbook of Pragmatics*. Blackwell Publishing.
- Kehler, A. (2011). Cohesion and coherence. In Heusinger, K. v., Maienborn, C., and Portner, P., editors, *Semantics: An International Handbook of Natural Language Meaning*, volume 2, pages 1963–1987. De Gruyter Mouton, Berlin.
- Kennedy, C. (2007). Vagueness and grammar: the semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy*.
- Kennedy, C. and McNally, L. (2005). Scale structure and the semantic typology of gradable predicates. *Language*, 81:345–381.
- Keshet, E. (2005). A matter of taste. Unpublished Manuscript, MIT.
- Kratzer, A. (1981). The notional category of modality. In Eikmeyer, H.-J. and Rieser, H., editors, *Words, Worlds, and Contexts: New Approaches in Word Semantics*, pages 38–74. De Gruyter, Berlin and New York.
- Kratzer, A. (1991). Modality. In von Stechow, A. and Wunderlich, D., editors, *Semantics: An International Handbook of Contemporary Research*, pages 639–650. de Gruyter, Berlin.

- Krifka, M., Pelletier, F. J., Carlson, G., ter Meulen, A., Link, G., and Chierchia, G. (1995). Genericity: An introduction. In Carlson, G. and Pelletier, F. J., editors, *The Generic Book*, pages 1 – 124. University of Chicago Press, Chicago.
- Kripke, S. (1980). *Naming and Necessity*. Basil Blackwell, Oxford.
- Lasnik, P. (2005). Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy*, 28(6):643–686.
- Lassiter, D. (2016). *Must*, knowledge, and (in)directness. *Natural Language Semantics*, 24(2):117–163.
- Lauer, S. (2013). *Towards a Dynamic Pragmatics*. PhD thesis, Stanford University.
- Lewis, D. (1973). Causation. *The Journal of Philosophy*, pages 556–567.
- López de Sa, D. (2008). Presuppositions of commonality: An indexical relativist account of disagreement. In García-Carpintero, M. and Kölbel, M., editors, *Relative Truth*, pages 297–310. Oxford University Press.
- Malamud, S. A. and Stephenson, T. (2015). Three ways to avoid commitments: Declarative force modifiers in the conversational scoreboard. *Journal of Semantics*, 32:275–311.
- Moltmann, F. (2010). Relative truth and the first person. *Philosophical Studies*, 150(2):187–220.
- Pearson, H. (2013). A judge-free semantics for predicates of personal taste. *Journal of Semantics*, 30(1):103–154.
- Peirce, C. S. (1877). The fixation of belief. *Popular Science Monthly*, 12:1–15.
- Piñon, C. (2006). Modal adverbs again. Presented at the Workshop on Syntax, Lexicon and Event Structure, The Hebrew University of Jerusalem.
- Portner, P. and Rubinstein, A. (2016). Extreme and non-extreme deontic modals. In Charlow, N. and Chrisman, M., editors, *Deontic Modality*. Oxford University Press, Oxford.
- Rawlins, K. (2010). Conversational backoff. In Li, N. and Lutz, D., editors, *Proceedings of SALT 20*, pages 347–365, Ithaca, NY. CLC Publications.
- Roelofsen, F. (2013). A bare bones attentive semantics for *might*. In Aloni, M., Franke, M., and Roelofsen, F., editors, *The dynamic, inquisitive, and visionary life of ϕ , $?\phi$, and $\diamond\phi$: a festschrift for Jeroen Groenendijk, Martin Stokhof, and Frank Veltman*, pages 190–215. Institute for Logic, Language, and Computation, Amsterdam.
- Rosenberg, J. (2012). P vs. NP Poll Results. *Communications of the ACM*, 55(5):10.
- Stalnaker, R. (1978). Assertion. In Cole, P., editor, *Syntax and Semantics 9: Pragmatics*. Academic Press, New York.
- Stalnaker, R. C. (2002). Common ground. *Linguistics and Philosophy*, 25:701–721.
- Stephenson, T. (2007). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30:487–525.

- Taranto, G. (2005). How discourse adjectives synchronize the speaker and the listener's beliefs. In Georgala, E. and Howell, J., editors, *Proceedings of SALT 15*, pages 241–257, Ithaca. Cornell University.
- Taranto, G. (2006). *Discourse Adjectives*. Routledge, New York.
- Vul, E., Goodman, N., Griffiths, T. L., and Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive science*, 38(4):599–637.
- Walker, R. (1996). A third parameter for unbound stress. In Kusumoto, K., editor, *Proceedings of North East Linguistic Society*, pages 441–456, Harvard University and MIT. Graduate Linguistic Student Association.
- Wolf, L. (2014). *Degrees of Assertion*. PhD thesis, Ben-Gurion University of the Negev, Beersheba.
- Wolf, L. and Cohen, A. (2011). Clarity as objectivized belief. In Égré, P. and Klinedinst, N., editors, *Vagueness and Language Use*, Palgrave Studies in Pragmatics, Language and Cognition Series, pages 165–190. Palgrave Macmillan, New York.
- Yalcin, S. (2011). Nonfactualism about epistemic modality. In Egan, A. and Weatherson, B., editors, *Epistemic Modality*, pages 295–332. Oxford University Press, New York.