

# Seeing Through the Clouds

John MacFarlane

## Abstract

Several recent authors have suggested that, in cases where a speaker's communicative intentions are not specific enough to determine a single proposition, we should think of a speaker as putting forward a *cloud* of propositions. This "putting forward" is supposed to be an illocutionary act, distinct from asserting each of the propositions in the cloud. How can we characterize the force of this speech act? What norms govern it? How does it affect the common ground? How does it function in communication? What is required for "uptake"? The standard stories about these things all presuppose that the content of an assertion is a single proposition, so any proponent of the cloudy picture owes us a new story. I argue that none of the proponents of proposition clouds has given an adequate answer to these questions. Instead of solving the problem of contextual indeterminacy by introducing novel speech acts whose contents are clouds of propositions, I propose, we should solve it by modifying our conception of propositions and adopting a form of expressivism.

Speech acts, we were all taught, have a force and a content. The force is what distinguishes assertions from questions, commands, and conjectures. The content is what distinguishes an assertion that it is raining from an assertion that it is snowing. Traditionally, the contents of assertions have been called *propositions*. In asserting that it is raining, I commit myself to the truth of the proposition that it is raining (at some particular time and place). If you want to agree or disagree with me, you can say *That is true* or *What he said is false*, where the terms *that* and *what he said* denote the proposition I asserted.

The orthodox picture of assertion goes along with an orthodox picture of communication. If I assert a proposition and you accept my assertion, you typically recognize what I have asserted and (unless you have reason to doubt me) acquire a belief whose content is the very same proposition. This proposition also gets added to a stock of propositions that are taken for granted in our conversation—the *common ground*. Typically the point of asserting a proposition to get others to accept it, to make it part of the common

ground. When your interlocutors do not recognize which proposition you meant to assert, your assertion has misfired: “uptake” has failed.

The orthodox story about assertion and communication is deeply embedded in the theoretical frameworks we use in semantics and pragmatics. But recently it has come under some pressure. Several theorists have proposed that the best way to make sense of assertions that are vague or indeterminate is to think of their contents as *clouds* of propositions, instead of single propositions (Braun and Sider 2007; Buchanan 2010; von Fintel and Gillies 2011). The idea has been invoked more recently in discussions of the metasemantics of contextual sensitivity (King 2014, 106) and contextualist views accounts of moral language (Khoo and Knobe 2018; Suikkanen, n.d.).

In what follows, I will take a critical look at this family of views. How should we understand the speech act we are making when we “put in play” a cloud of propositions? How exactly must the orthodox picture of speech acts and communication be modified to make room for this idea? I find different answers to these questions in Braun and Sider (2007), Buchanan (2010), and von Fintel and Gillies (2011), but none of the answers yield a plausible account that can rival the orthodox one. Although vagueness and contextual underdetermination do pose a problem for the orthodox conception of assertion and communication, clouds of propositions are not the right response to this problem. In closing, I will sketch what I think is a better approach.

## 1 MOTIVATING CLOUDS

Although the three versions of the proposition-cloud view differ in significant ways and focus on different bits of language, they are motivated by similar considerations. The authors present cases in which the speaker’s intentions do not single out one of a number of possible assertible contents. They then propose that, instead of taking the speaker to be expressing a single proposition, we take her to be expressing a whole cloud of them.

Braun and Sider focus on vague expressions. They assume that the world itself is not vague, and that propositional truth is not relative to anything other than the state of the world.<sup>1</sup> It follows that every proposition is either determinately true or determinately false at a given world. That is, propositions are precise. Vague sentences, then, do not express propositions, even in context. Neither the meanings of vague expressions nor

---

<sup>1</sup>One might object to these assumptions, of course. But our aim here is to explore the “cloud of propositions” approach, not to consider every possible approach to vagueness.

anything about the context (including speakers' intentions) are enough to pin down a single proposition. Even in a specific context, a vague sentence expresses a cloud of propositions. For example, the sentence *Bill is tall* might express a cloud containing the proposition that Bill exceeds 196 mm in height, the proposition that Bill exceeds 197 mm in height, and many others. As Braun and Sider put it:

There is typically a cloud of propositions in the neighborhood of a sentence uttered by a vague speaker. Vagueness prevents the speaker from singling out one of these propositions uniquely, but does not banish the cloud. (Braun and Sider 2007, 135)

Buchanan's appeal to clouds of propositions is motivated not by lexical vagueness, but by indeterminacy in the resolution of contextual sensitivity. For example, suppose Chet asserts

(1) Every beer is in the bucket.

Which proposition he has asserted depends on how the quantifier domain is specified in context. Does Chet mean *every beer we bought at the bodega*, *every beer we will serve at the party*, *every beer for our guests*, or *every beer at the apartment*? Buchanan suggests, plausibly, that Chet's intentions may not distinguish between these options. If asked to clarify what he meant, he might respond with one of these propositions. But there is no reason to think that he had this one, rather than the others, in mind when he uttered (1).

The fact that the speaker might, as it were, "fall back" on any one, or more, of the foregoing candidates, suggests that no single such candidate, or set of candidates, perfectly capture his communicative intentions in uttering [(1)]. Chet's communicative intentions, such as they are, exhibit a certain kind of *generality* and *indifference* that precludes us from identifying any one of the candidate propositions as *the one he meant*. (Buchanan 2010, 350)

Buchanan concludes that

The object of Chet's communicative intentions is not a proposition, but rather a property of propositions. ... Chet's utterance is, in some sense, "associated" with many non-truth conditionally equivalent propositions—namely, those propositions that are of the intended (restricted) type. (Buchanan 2010, 358)

King (2014, n.d.) gives many more examples of this kind of "felicitous underdetermination." For example, looking a group of surfers down the beach, one might say

(2) Those guys are good.

Which plurality of surfers in that area is being referred to? The speaker may have no precise idea. Planning a wedding in California, one might say

(3) Let's go with a local firm.

Local to the town, the county, the larger area? Nothing about the speaker's intention may decide that. King takes the vagueness of gradable adjectives (which motivated Braun and Sider) to be a special case of the more general phenomenon of felicitous underdetermination. In degree semantics for gradable adjectives (e.g., Kennedy 2007), the extensions of gradable adjectives are relativized to a contextually provided threshold. But when one says

(4) The ocean was cold today,

one doesn't generally have in mind a particular threshold: a particular temperature below which the ocean counts as cold. King comments that "it is natural to think that in such cases speakers' intentions determine a range of degrees on the relevant scales" (King 2014, 112). Although King does not defend any particular view about how to understand assertions in the presence of contextual underdetermination, he mentions the cloud-of-propositions view approvingly (King 2014, 106).

von Fintel and Gillies also motivate the idea of a cloud of propositions using a case of felicitous underdetermination, which they attribute to Chris Potts:

Billy meets Alex at a conference, and asks her:

(18) Where are you from?

That question is supposed, given a context, to partition answer-space according to how low-level in that context Billy wants his details about Alex to be. But notice that it's not really clear whether Billy wants to know where Alex is currently on sabbatical or where Alex teaches or where Alex went to graduate school or where Alex grew up. And—the point for us—Billy might not know what he wants to know. He just wants to know a bit more about Alex and will decide after she answers whether he got an answer to his question or not. He doesn't have to have the level of granularity sorted out before he asks the question. So context (or context plus Billy's intentions) need not resolve the contextual ambiguity. (von Fintel and Gillies 2011, 118)

To handle such cases, they suggest,

we can think of utterances taking place against a cloud of admissible contexts. ...

There is no such thing as 'the context', only the contexts admissible or compatible with the facts as they are. (von Fintel and Gillies 2011, 118)

They apply this general idea to epistemic modals, which are standardly taken to be sensitive to a contextually supplied body of information. Recent critics of standard contextualist views have argued that in many cases *no* setting for this contextual parameter can explain both the readiness of speakers to make epistemic possibility assertions and the readiness of listeners to reject these assertions.<sup>2</sup> For example, Alex might assert

(5) The keys might be in the car

even when, for all Alex knows, Billy may know that the keys aren't there.<sup>3</sup> To explain this, it looks as if we need to restrict the contextually relevant information to what Alex knows. But this yields the wrong predictions about the significance of agreeing or disagreeing with Alex's assertion. If Billy agrees with Alex, saying

(6) That's right, they might be,

she will normally be indicating that it is consistent with *her* (Billy's) information that the keys are in the car—not that it is consistent with Alex's information. If Billy knows that the keys aren't in the car, then it is inappropriate for her to agree with Alex:

(7) # That's right, but I know they're on the table.

Instead, she ought to disagree:

(8) No, they can't be in the car: I just checked there.

So agreeing or disagreeing with Alex's assertion of (5) is not agreeing or disagreeing that it is consistent with what Alex knows that the keys are in the car. To make sense of this, it seems we must take the information relevant to Alex's assertion of (5) to include what Billy knows. But then it becomes hard to understand how Alex could have been warranted in asserting (5) in the first place. (And the difficulty increases if we consider eavesdroppers not known to the speaker.) There seems to be no setting of the contextually relevant information that can explain both Alex's entitlement to assert (5) and Billy's entitlement to agree or disagree.

In response, von Fintel and Gillies argue that there is a kind of contextual indeterminacy here. The speech situation is compatible with a cloud of possible settings for the relevant information state. Hence Alex's assertion cannot be identified with any *one* of the

---

<sup>2</sup>See Price (1983) (considering probability rather than possibility); MacFarlane (2003); MacFarlane (2011); MacFarlane (2016 ch. 10); Egan, Hawthorne, and Weatherston (2005); Yalcin (2007).

<sup>3</sup>This is von Fintel and Gillies' example, but it mirrors similar examples used by the relativist and expressivist critics of contextualism cited above.

propositions one gets by plugging a particular body of information into the schema

(9) It is compatible with information  $\mathcal{F}$  that the keys are in the car.

Instead, in uttering (5) Alex expresses a *cloud* of propositions: the proposition that it is compatible with Alex's information (A) that the keys are in the car, the proposition that it is compatible with Billy's information (B) that the keys are in the car, and the proposition that it is compatible with their combined information (AB) that the keys are in the car. Alex is warranted in making her claim because she has warrant for one of the propositions in the cloud (A). But Billy can pick one of the other propositions (AB) to respond to.

Though von Fintel and Gillies are mainly concerned with the case of epistemic modals, the way they motivate the cloud-of-propositions picture suggests that they take it to be the appropriate response to contextual underdetermination in general. The case they use to introduce the picture has nothing to do with modals. And they concede that their strategy would seem ad hoc if the cloud of propositions view didn't apply also in other cases, such as implicit quantifier domains (von Fintel and Gillies 2011, 123).

## 2 HOW TO DO THINGS WITH CLOUDS

It is natural enough to say that, in cases of contextual underdetermination, a speaker expresses a cloud of propositions. But in making a speech act, one does more than express contents. One asserts them, supposes them, asks about them. In order to understand the cloudy picture, then, we need to understand what it is to make an assertion (or other speech act) using a cloud of propositions. Here, as we shall see, the three views we are considering give very different answers.

Right away we face a terminological issue. The term *proposition* is sometimes stipulated to mean the content of assertion (Cartwright 1962). On this way of talking, the thesis that what is asserted is a cloud of propositions is nonsensical. It may turn out that what is asserted is a cloud of somethings, but if so, the cloud itself is the content of the assertion, and hence a proposition, and the somethings that compose it are something else.

There are two ways forward if we want to make sense of the cloudy picture: we can say that the speech acts one is making when one deploys a cloud of propositions are not assertions, or we can reject the idea that propositions are the contents of assertions.

von Fintel and Gillies take the first approach. They talk of speakers "putting in play" a

cloud of propositions in situations where they wouldn't be entitled to "flat-out assert" all of them (von Fintel and Gillies 2011, 119–20). That looks like a denial that the speech act is one of assertion (see also 2011, 117 n. 18). The use of the terminology "putting in play" is unfortunate. Intuitively, one can put propositions in play by asking a question, making a conjecture, or performing many other kinds of illocutionary acts, but von Fintel and Gillies are using the term for a specific kind of speech act, whose force they will go on to describe. It might be better to introduce a new term, like *cloudserting*. But as long as we keep firmly in mind that *putting in play* is supposed to be a specific, assertion-like speech act whose content is a cloud of propositions, we can avoid confusion.

Braun and Sider take the second approach: they say that assertions can have clouds of propositions as their contents. This commits them to rejecting the idea that propositions are the contents of assertions.

Buchanan is more difficult to pin down on these issues. He avoids the terminology of *assertion* and *speech act* entirely, talking instead of what the speaker means, so he does not take a stand on whether the speech acts in question are assertions. He says that propositions are the "objects of belief and certain other of our cognitive attitudes." He also assumes that propositions "determine truth-conditions" (Buchanan 2010, 341). It seems open to him to say either that some assertions have clouds of propositions as their contents, or to say that the acts of speaker-meaning in question are not, strictly speaking, assertions.

It is important not to lose sight of these terminological issues, but I don't think it matters much how these terminological issues are resolved, because the substantive issue can be raised either way. What is important is to see that there *is* a substantive issue. On all of the views we are considering, we can use indicative sentences to "put forward" clouds of propositions. How can we characterize the force of this speech act? What norms govern it? How does it affect the common ground? How does it function in communication? What is required for "uptake"? The standard stories about these things all presuppose that the content of an assertion is a single proposition, so any proponent of the cloudy picture owes us a new story.

### 3 BRAUN AND SIDER

One natural view is that putting forward a cloud of propositions is committing oneself to the truth of every proposition in the cloud. That is essentially the view of Braun and

Sider (2007).

On Braun and Sider's view, vague sentences lack truth values in context, even when they do not involve borderline cases. For a sentence to be true in a context, they think, there must be a unique proposition that it expresses at that context, and that proposition must be true. The uniqueness condition fails for sentences containing vague expressions, even for non-borderline sentences like

(10) A person with no hair is bald.

So all such sentences lack truth values. Truth, then, "is an impossible standard that we never achieve" (Braun and Sider 2007, 135). However,

...it is usually harmless to *ignore* vagueness, set it aside, and act as if one's sentence is not vague, but rather expresses a unique proposition. When vagueness is being ignored, the cooperative communicator satisfies her communicative obligations well enough by uttering sentences that are approximately true... (Braun and Sider 2007, 135)

To say that a vague sentence is *approximately true* is to say that all of the propositions associated with it—the propositions that would be expressed by it on various legitimate disambiguations of its vague expressions—are true. For example, (10) is approximately true, because on any legitimate way of disambiguating *bald* it will express a true proposition.

Thus, when we are ignoring vagueness (as we usually are in everyday life), the norm for asserting or "putting forward" a cloud of propositions is that each proposition in the cloud be true. Asserting a cloud is committing oneself to the truth of every proposition in the cloud. And the effect of accepting such an assertion is presumably the addition of all of these propositions to the common ground. (When we aren't ignoring vagueness, on the other hand, no assertion of a cloud of propositions can meet the standard for assertion.)

This is a nice, simple story, and it is compatible with orthodox views of propositions and of the common ground. But I do not think it is a plausible story. It makes the condition for making vague assertions too stringent. As Barker (2002, 2009, 2013) has noted, a speaker might well assert

(11) Richard is tall

in order to further constrain the range of legitimate disambiguations of *tall* in that context. In such a case, one might admit that there are (prior to one's assertion) legitimate



ways of precisifying *tall* on which Richard would not count as tall. The force of the assertion is to propose that we go on in such a way that, within our conversation at least, Richard does count as tall.

Though Braun and Sider do not consider this objection, it is one they might meet by appealing to *accommodation* (Lewis 1979). Accommodation, as Lewis describes it, is the process by which nonfactual contextual parameters that affect the interpretation of utterance—which he conceives as “components of conversational score”—are adjusted when needed in order to interpret an utterance as making a reasonable move in the conversation. For example, suppose the conversation is currently governed by a high “standard of precision,” so that geometrical terms like *hexagonal* apply only to figures with very straight sides, and someone asserts

(12) France is hexagonal.

The proposition (12) expresses given the current standard of precision is that France has six perfectly straight sides, and this is already ruled out by the common ground. So, in order to interpret the assertion as one that might be true, we will take the standard of precision governing the conversation to have tacitly relaxed.

By appealing to this mechanism, Braun and Sider could explain how an assertion of (11), in a borderline case, can have the effect of contracting the range of legitimate disambiguations of *tall*. If there are legitimate disambiguations on which Richard does not count as tall, then (11) cannot meet the standard of approximate truth. So the participants in the conversation will naturally accommodate the speaker, adjusting the range of legitimate disambiguations on the conversational scoreboard so that Richard counts as tall on every legitimate disambiguation.

However, this reply only works in the case where Richard’s height is mutually known. For only in that case do we know how to adjust the range of legitimate disambiguations of *tall* so that Richard counts as *tall* on all of them. Let’s imagine, then, that Richard is not in the room for us to see. We’ve all seen him before, but we’re unsure about his exact height: as far as the group knows, it could be anywhere between 190 and 195 cm. Suppose the prior range of legitimate disambiguations for *tall* allows thresholds from 185 to 195 cm. What does our story say, now, about the force of asserting (11)?

In this scenario, the current range of legitimate disambiguations does not preclude the assertion’s being approximately true. For if Richard is 195 cm—which is an open possibility given the common ground—then he exceeds the threshold for *tall* on any legitimate

way of disambiguating *tall*. So accommodation is not triggered in this case. The update proposed by the assertion of (11) is that Richard is 195 cm tall—or so the framework predicts.

But this prediction is wrong. To see why, suppose that we find out that Richard is 191 cm tall. According to the theory we are considering, the common ground already excludes this possibility, so the context “crashes” and must be repaired. In addition, the earlier assertion of (11) must be regarded as false and retracted, just as an earlier assertion of

(13) Richard is 195 cm tall

would have to be. But in fact, we can learn that Richard is 191 cm tall without retracting (11), and without a crash in the context. Once we learn Richard’s height, the earlier assertion of (11) commits us to changing the range of legitimate disambiguations for *tall* to the range 185 cm to 191 cm, so that Richard counts as tall no matter how *tall* is disambiguated. (11) is, then, a kind of conditional commitment: a resolution to contract the range of disambiguations for *tall* as needed, when new information about Richard’s height comes in. It is not, as Braun and Sider’s theory would predict, an unconditional commitment about Richard’s height.

Nor does this problem come from some idiosyncratic quirk of Braun and Sider’s account. No theory that models the common ground as a set of worlds plus a range of legitimate disambiguations can represent this kind of conditional update. Since asserting (11) doesn’t require any specific update to the range of legitimate disambiguations for *tall*, we can only think of it as an update to the other component of the common ground—the set of worlds. (In the final section we will present a model of common ground that fixes this problem.)

Even if Braun and Sider’s proposal worked for vague predicates, it would not be useful for the other cases of felicitous contextual underdetermination considered by Buchanan, King, von Stechow and Gillies. Consider a variant of (2): we look down the beach and say, of a group of surfers with indeterminate boundaries,

(14) Those guys are the only ones on the beach who really know how to surf.

Suppose that *those guys* has two legitimate interpretations:

- (a) *those guys* = Abe, Bob, Cindy, Maria, Zeke
- (b) *those guys* = Abe, Bob, Cindy, Maria, Sid

If we were committing ourselves to (14) being true on every legitimate disambiguation

of *those guys*, then the commitment would be inconsistent. For (14) to be true on interpretation (a), it must not be the case that Sid really knows how to surf. But for it to be true on interpretation (b), it must be the case that Sid really knows how to surf.

Or consider von Fintel and Gillies' case of Alex, Billy, and the keys. If assertoric commitment is commitment to the truth of all the propositions in the cloud, then in asserting that the keys might be in the car, Alex is committing herself to the proposition that Billy's information does not rule out the key's being in the car. But that seems too strong. Alex isn't in a position to make a claim about what Billy's information leaves open; she is asserting (5) partly to flush out any relevant information Billy might have.

We need a different story, then, about the force of asserting or otherwise "putting in play" a cloud of propositions.

#### 4 BUCHANAN

Buchanan rejects the view that when a speaker asserts a cloud of propositions (or, as he prefers to say, a *proposition type*), the intended update is to add the conjunction of the propositions to the common ground. The conjunction, he notes, "is simply another candidate proposition that Chet did not mean" (Buchanan 2010, 353). This is shown, he thinks, by the fact that Chet's audience "need not entertain each of [the propositions in the cloud] in order to understand the utterance" (Buchanan 2010, 353). It is sufficient for uptake that the audience "entertain *any one*, or more, of the candidates on the basis of the utterance," and be "thereby *disposed* to accept some number of the other salient candidates" (Buchanan 2010, 366, n. 22). Thus,

*understanding a speaker's utterance requires entertaining some one or more propositions which are of the restricted proposition-type the speaker meant.* In the case of the utterance of [*Every beer is in the bucket*], Tim need not entertain the restricted proposition-type that Chet means; rather, what is required is that Tim entertain one or more of the candidates of that type on the basis of Chet's utterance. (Buchanan 2010, 359)

To understand Chet's assertion of (1), then, Tim need not grasp the proposition-type or cloud Chet intended; he need only fasten on one of the propositions in this cloud. It doesn't matter which.

Buchanan isn't specific here about how he sees the proposed update to the common ground, but what he says makes it clear that grasping the proposed update cannot re-

quire recognizing what proposition type (or cloud of propositions) the speaker intended. I might fasten on one proposition,  $P_1$ , while you fasten on another,  $P_2$ . This is, for Buchanan, a *feature* of the view: he wants to explain how “an utterance might be understood in non-equivalent, yet equally correct ways” (Buchanan 2010, 359). But this feature makes it impossible to see how the speech act could be viewed as proposing any specific update to a *common* ground. Suppose Chet asserts (1), and Tim and Zeke both signal their assent. What has been added to the common ground? It may be that Tim has fastened on the proposition that every beer they bought in the bodega is in the bucket, while Zeke has fastened on the proposition that every beer they will serve at the party is in the bucket. These are different updates to the common ground.<sup>4</sup> If there’s no common knowledge which of these updates has been accepted, then we can’t think of them as adding to the common ground.

Buchanan presents himself as rejecting what might seem a peripheral part of Grice’s account of speaker meaning—the idea that what is communicated is a proposition—while keeping the basic shape of the account. But his view commits him to rejecting the *central* idea of Grice’s account: the idea that the speaker’s meaning intention is *inter alia* an intention that the audience recognize this intention.<sup>5</sup> It is this essential transparency, Grice thinks, that distinguishes meaning intentions from other kinds of intentions to produce effects on hearers. And, as Stalnaker (2002, 74) observes, it is the transparency of Gricean meaning intentions that allows us to think of speech acts as updating a common ground. Buchanan rejects the transparency of meaning intentions. On his view, the speaker’s intention is satisfied if the audience entertains *some* proposition of the relevant type. But the speaker need not intend that the audience recognize this intention, for that would require recognizing the proposition type the speaker intended, and Buchanan denies that this is necessary for successful uptake.

Would it help to give up this part of Buchanan’s view, and say that uptake requires recognizing the proposition type intended by the speaker?<sup>6</sup> A natural thought is that the speaker intends the hearer to recognize the type intended, but allows the hearer to pick *which* token of that type is to be added to the common ground. This would allow us to

---

<sup>4</sup>Unless it is already common ground that they will only serve the beers they bought at the bodega.

<sup>5</sup>Buchanan seems to recognize that he needs to modify this part of Grice’s view (see Buchanan 2010, 368 n. 34).

<sup>6</sup>I assume that Buchanan resists this because many of the same worries about determinacy that arose for propositions can be raised for proposition types. That is, there are many proposition types—and accordingly many distinct clouds of propositions—that a speaker might be taken to have asserted. But let us leave this issue aside for now, since our task is to see what sense might be made of asserting a cloud of propositions.

continue to model the common ground as a set of accepted propositions—but only if it becomes mutually known which proposition the hearer has fastened on to. This would happen if Tim responds to Chet’s assertion of (1) by saying,

(15) Yes, every beer we bought at the bodega is in the bucket full of ice on the back porch,

thus clarifying the update. But what about the more normal case where Tim simply responds with *Yep* or a nod? Then there will be no common knowledge about what the update is supposed to be—and that is just to say that we cannot view the update as an update to a *common* ground.

If there were some algorithm for determining which proposition the hearer was assenting to, then an account like this could succeed. But of course, there isn’t such an algorithm. A hearer could be picking up on any of the propositions in the cloud. As we will see, von Fintel and Gillies’ proposal can be seen as an attempt to solve this problem.

## 5 VON FINTEL AND GILLIES

The view sketched by von Fintel and Gillies is designed to give clear predictions about how the common ground should be updated after various kinds of response to an epistemic possibility claim. Unfortunately, as I’ll argue, although it relies on principles that should be generally applicable if they are valid at all, it only gives good results in the specific case of epistemic possibility claims.

Recall that on von Fintel and Gillies’ view, a bare epistemic modal claim like

(16) It might be that  $P$

“puts in play” a cloud of propositions of the form

(17) It is compatible with information  $\mathcal{F}$  that  $P$ ,

with values for  $\mathcal{F}$  taken from a contextually determined range. We have been asking what it is to “put in play” a cloud of propositions of this kind. What norms govern this speech act, and how does it affect the common ground?

In answer to this question, von Fintel and Gillies articulate two norms: one governing the making of this speech act and another governing its uptake (acceptance or rejection). The norm for making the speech act is

ASSERT One may put in play a cloud of propositions just in case one is in a position to “flat out assert” one of the propositions in the cloud. (von Fintel and Gillies 2011, 120)

Two exegetical notes, before we continue. First, although von Fintel and Gillies state ASSERT as a principle governing utterances of sentences of the form *It might be that  $\phi$* , I am taking it as a more general pragmatic principle, since our interest is in exploring what sense, in general, can be made of “putting in play a cloud of propositions.” If ASSERT turns out to be plausible only for epistemic possibility claims, then one would want to look for a more generally applicable principle from which it follows, given special features of epistemic possibility modals. Otherwise ASSERT looks ad hoc.

Second, although I have formulated ASSERT as a necessary and sufficient condition, von Fintel and Gillies’ wording suggests only a necessary condition:

Suppose an utterance of *might(B)( $\phi$ )* by *S* puts in play the propositions  $P_1, P_2, \dots$ . Then *S* must have been in a position to flat out assert one of the  $P_i$ ’s.

Our proposal is that in order for a speaker to be within her linguistic and epistemic rights when she issues a BEM [bare epistemic modal sentence] against a cloud of contexts, she has to be in a position to flat out assert one of the meanings it can have, given that cloud. (von Fintel and Gillies 2011, 120)

However, that their argument requires not just a necessary but a sufficient condition for the permissibility of an assertion. Applying ASSERT to their scenario, they say:

given the facts of that scenario, Alex is justified in uttering the BEM *iff* she is justified in claiming that her evidence does not rule out the prejacent. (von Fintel and Gillies 2011, 120, emphasis added)

They are arguing as follows: Alex is entitled to flat-out assert that, given what *she* knows, the keys might be in the car. This is one of the propositions in the cloud associated with (5): the one they call the “*A*-reading.” Therefore, Alex is entitled to put in play the cloud. For this to follow, though, we need a *sufficient* condition for the permissibility of putting in play the cloud, not just a necessary condition. Accordingly, I have presented ASSERT as a necessary and sufficient condition.

ASSERT gives us one piece of the puzzle: it shows why Alex can be entitled to assert (5), even though she doesn’t know what Billy knows about the location of the keys. What about the other piece? Why is Billy entitled to reject Alex’s claim, if she knows the keys aren’t in the car?

To sort this out, we need to understand what the hearer is supposed to do when a speaker “puts in play” a cloud of propositions. The basic idea is that the hearer selects one of the propositions in the cloud to react to, accepting or rejecting it. Thus, by putting in play a cloud of propositions, instead of a single one, the speaker cedes some control to the hearer in determining what update to the common ground is to be made. If the hearer selects the proposition  $P$ , and accepts it, then  $P$  is added to the common ground. If she selects  $P$  and rejects it, then the negation of  $P$  is added to the common ground.

For this story to work, though, it must be possible for all parties to the conversation to figure out which proposition has been targeted by the hearer. Otherwise, there won't be any common understanding about how the common ground is to be updated when the hearer accepts or rejects the speaker's claim. So, the hearer can't be given complete freedom to target any proposition from the cloud:

...not just any one of them will do. Instead, we argue that the hearer is guided by what response to which proposition will be most informative in the conversation. When the modal is an existential like *might*, this will in fact lead to a dominance of negative replies. (von Fintel and Gillies 2011, 121)

von Fintel and Gillies articulate this constraint on the hearer's choices through a principle they call

CONFIRM/DENY One may confirm (deny) a speech act that puts in play a cloud of propositions just in case one takes the strongest proposition in the cloud that one reasonably has an opinion about to be true (false).<sup>7</sup> (von Fintel and Gillies 2011, 121)

Let's see how this applies to the case of Alex and Billy. In asserting (5), Alex puts in play a cloud containing three propositions:

- $\text{Poss}_A K$ : Alex's information doesn't exclude the keys being in the car.
- $\text{Poss}_B K$ : Billy's information doesn't exclude the keys being in the car.
- $\text{Poss}_{AB} K$ : Alex and Billy's information pooled together doesn't exclude the keys being in the car.

$\text{Poss}_{AB} K$  is logically stronger than  $\text{Poss}_A K$ : if Alex and Billy's combined information is compatible with the keys being in the car, then Alex's information also must be compatible. So CONFIRM/DENY tells Billy to target  $\text{Poss}_{AB} K$ , provided she has a reasonably

---

<sup>7</sup>As with ASSERT, I have restated this as a general principle about cloudy speech acts, rather than bare epistemic modal claims.

grounded opinion about its truth. Since Alex and Billy mutually know, in this case, that Billy is in a position to have an opinion about the truth of  $\text{Poss}_{AB} K$ , it becomes mutually known that this is the proposition Billy was targeting, and the common ground can be updated accordingly.

There is a puzzle here about how CONFIRM/DENY is supposed to relate to the injunction of making the most informative response. After all, accepting a stronger proposition is more informative than accepting a weaker one, but *rejecting* the stronger proposition is less informative than rejecting the weaker one. Yet CONFIRM/DENY says that the hearer should target the strongest proposition, whether she is confirming or denying it. According to von Fintel and Gillies, the most informative move Billy can make is to reject  $\text{Poss}_{AB} K$ . But presumably Billy is also in a position to reject  $\text{Poss}_B K$ , which is weaker than  $\text{Poss}_{AB} K$ , and rejecting a weaker proposition would be more informative than rejecting a stronger one. If we stick to the “most informative move” motivation, then, we should predict that if Billy *accepts* Alex’s claim, he is targeting  $\text{Poss}_{AB} K$ , while if he *rejects* it, he is targeting  $\text{Poss}_B K$ .

If we are going to take into account all of the possible responses Billy might make, we also need to be able to say whether rejecting the stronger proposition  $\text{Poss}_{AB} K$  would be more informative than *accepting* the weaker proposition  $\text{Poss}_A K$ . In the scenario as described, accepting  $\text{Poss}_A K$  is not very informative: it is probably already common ground after Alex’s utterance that she doesn’t know where the keys are. So, in this particular case, rejecting  $\text{Poss}_{AB} K$  is more informative than accepting  $\text{Poss}_A K$ . But that does not follow from the logical relation between  $\text{Poss}_{AB} K$  and  $\text{Poss}_A K$ ; it depends on special features of the case. There are certainly cases in accepting a logically weaker proposition would be more informative than rejecting a logically stronger one. For example, if the question at hand is who is be coming to the party, then accepting the proposition that Sarah is coming is more informative than rejecting the (stronger) proposition that Sarah and the entire population of Des Moines, Iowa are coming.

The upshot is that the instruction given in CONFIRM/DENY—always target the strongest proposition in the cloud about whose truth value you have a reasonable opinion—conflicts with the intuitive motivation of maximizing the informativeness of one’s contribution. I am not sure whether to give precedence to the ideal of maximizing informativeness or the explicit statement of CONFIRM/DENY, so in what follows I will consider both options.

What is attractive about von Fintel and Gillies’ proposal is that it offers a way of determin-



ing which proposition the hearer is fastening on, and hence how the common ground should be updated. Thus, if Billy rejects Alex's claim, the common ground can be updated with  $\neg \text{Poss}_{AB} K$  (von Fintel and Gillies 2011, 123–4). On the other hand, if Billy accepts Alex's claim, then it becomes common ground that  $\text{Poss}_{AB} K$  is true. In either case, the exchange leads to a standard update of the common ground, conceived in the conventional way as a set of worlds:

...once a hearer has confirmed or denied the BEM with all its indeterminacy, the resulting common ground is quite determinate. (von Fintel and Gillies 2011, 124)

It is CONFIRM/DENY (or perhaps the ideal of maximizing informativeness) that allows the parties to reason in this way. Without this instruction for determining which proposition the hearer is targeting, von Fintel and Gillies' story would be much like Buchanan's, and it would suffer from the same flaw: it would be unclear how the common ground is to be updated after a claim is accepted or rejected.

The problem is that this view does not generalize well beyond the case to which von Fintel and Gillies apply it: bare epistemic possibility claims. The picture requires that the hearer can figure out in a transparent way which of the propositions in the cloud the hearer is accepting or rejecting. Otherwise we wouldn't have a well-defined update to the common ground. In the case of a bare epistemic *might* statement, von Fintel and Gillies claim, that is going to be the strongest proposition in the cloud that the hearer "reasonably has an opinion about." In our toy example, it's clear to both parties that Billy is in a position to have a reasonable opinion about (5). So Alex can come to know which proposition Billy is rejecting, and it is transparent how the common ground is to be updated.

But now consider a different sentence:

(18) The keys are probably in the car.

Our cloud now consists of these propositions:

- $\text{Prob}_A K$ : It is probable given Alex's information that the keys are in the car.
- $\text{Prob}_B K$ : It is probable given Billy's information that the keys are in the car.
- $\text{Prob}_{AB} K$ : It is probable given Alex and Billy's information pooled together that the keys are in the car.

Suppose that  $\text{Prob}_A K$  is true. Then, according to ASSERT, Alex is warranted in asserting (18). Now what about Billy? Should she accept or reject Alex's statement? Ac-

ording to CONFIRM/DENY, she should identify the strongest proposition in the cloud whose truth value she has a reasonable opinion about, and target that one. But in this case, none of our three propositions entails any of the others, so there isn't a strongest.<sup>8</sup> If she *does* agree with or reject Alex's claim, then, it will not be clear how to update the common ground.

Thus it looks as if CONFIRM/DENY can only do the job it is meant to do—showing us how the common ground is to be updated after speech act that “puts in play” a cloud of propositions is accepted or rejected—in the core case of epistemic possibility modals, and it only works there because of a special property of these modals (monotonicity) that is not even shared by epistemic modals in general.

In this particular case, one might try to salvage things by (a) arguing that Billy is not in a position to have a reasonable opinion about  $\text{Prob}_{AB} K$ , and (b) forgetting about CONFIRM/DENY and reasoning informally about informativeness, instead of focusing on logical strength. Suppose it is common knowledge that the only propositions in the cloud Billy is in a position to have a warranted opinion about are  $\text{Prob}_A K$  and  $\text{Prob}_B K$ . Since Alex already knows  $\text{Prob}_A K$  but doesn't know the truth value of  $\text{Prob}_B K$ , it is more informative for Billy to target  $\text{Prob}_B K$ . So, as long as Alex knows all of this, and knows that Billy doesn't have an opinion on  $\text{Prob}_{AB} K$ , she can work out how the common ground is to be updated if Billy accepts or rejects her claim.

However, even if Billy is not in a position to have a reasonable opinion about  $\text{Prob}_{AB} K$ , it is hard to see how this could be common ground between Alex and Billy. After all, for all Alex knows, it may be that Billy's information is strictly stronger than Alex's. That is, for all Alex knows, it may be that Billy knows everything that Alex knows that is relevant to the location of the keys, plus more in addition, and it may be that Billy knows this fact. In that case, something would be probable given  $AB$  just in case it is probable given  $B$ , so Billy *would* be in a position to have a reasonable belief about  $\text{Prob}_{AB} K$ .

Nor do we have to go far to find other cases with multiple possible updates, none of which is strictly more informative than any of the others. Consider our variant (14) of King's case of the surfers on the beach. In this case there are two possible interpretations,

---

<sup>8</sup>This is because Poss is monotonic in a way that Prob is not. Adding information can only remove possibilities, not add them, so if  $K$  is possible given  $AB$ , it must be possible given  $A$ . By contrast, adding information might, depending on the case, make something either more probable or less probable. (Simple example: learning that a horse has won all its previous races will increase our subjective probability that it will win this race; learning that it has lost all its previous races will decrease it.) So we have no entailment either way between  $\text{Prob}_A K$  and  $\text{Prob}_{AB} K$ .

- (a) *those guys* = Abe, Bob, Cindy, Maria, Zeke
- (b) *those guys* = Abe, Bob, Cindy, Maria, Sid.

Neither of these interpretations entails the other. Nor is there any looser sense in which one would be a more informative contribution to the conversation than the other.

## 6 THE DEEPER PROBLEM

One might try to salvage the core of the cloudy view while throwing out the idea that there is an algorithm (such as CONFIRM/DENY) for determining which proposition the hearer is targeting. Instead, one might say that the hearer has an obligation, in accepting or rejecting the speaker's claim, to make it clear to the speaker which proposition she is targeting. She might do this directly:

- (19) No, it isn't probable given what we both know that the keys are in the car.

Or might do it indirectly—for example, by backing up her claim with evidence that would only be relevant to one of the propositions in the cloud:

- (20) No, it isn't likely to be in the car; I am pretty sure that I would have noticed it when I was in the garage just now.

This would preserve the idea that the update is in some way a matter for negotiation between the speaker and hearer, while giving up the idea that both parties can figure out which proposition the hearer is targeting without explicit hints.

But I don't think this will help, in general, with the cases of felicitous underdetermination that have typically motivated appeals to clouds. The deeper problem is that, when hearers accept or reject, they are typically *not* targeting a single, determinate proposition from the cloud, any more than speakers are.

Consider Buchanan's case (1). It's plausible that the speaker's intentions are compatible with a cloud of propositions that differ in exactly how the incomplete definite descriptions *the beer* and *the bucket* are supplemented. But the same goes for the *hearer's* intentions. Chet asserts (1); Tim assents with a laconic *Yep*. If we now ask Tim whether he was assenting to

- (21) Every beer we bought at the bodega is in the bucket full of ice on the back porch,  
or

(22) Every beer we will serve at the party is in the bucket on the back porch, or

(23) Every beer in the apartment is in the bucket with pictures of pirates on it,

will he be able to answer? Even if he does answer, does his answer reflect a determinate intention he had when he assented to (1), or has he just plumped for something retrospectively? In this case, the indeterminacy left open by the speaker doesn't seem to be resolved in the hearer's response.

Or consider King's case of the surfers down the beach. Suppose King says (2) and his hearers reject his claim, saying,

(24) No, they're not that good, they just got lucky on that wave.

Do his hearers need to have any more determinate an idea than King did of who, exactly, belongs to the plurality denoted by *they*?

The problem is perhaps clearest in the case of vague gradable adjectives. Suppose Anna sees Tim, who stands 190 cm tall, and says,

(25) He's tall,

thereby putting in play a cloud of propositions of the form

(26)  $T_n$  = Tim is at least  $n$  mm tall,

for every  $n$  between 1850 and 2000. She is entitled to do this, according to *ASSERT*, as long as she knows that Tim's height is at least 1850 mm, for that is the condition for the weakest of these propositions to be true. It would be completely natural for her hearers to accept or reject Anna's claim. But there is no plausibility to the idea that, in doing so, they are singling out a specific proposition in the cloud. Their intentions are just as indeterminate as Anna's. Nor does *CONFIRM/DENY* (or the informal guidance to maximize informativeness) give any help here, even though here, the propositions in the cloud stand in clear relations of logical strength. Suppose Anna's hearers accept her claim. We are supposed to update the common ground with the strongest proposition in the cloud they have a reasonable opinion is true. Presumably it is reasonable for them to believe that Tim is at least 1850 mm tall, since that is already common ground. What about 1851 mm? 1852 mm? Even if there is a fact of the matter about what is the greatest  $n$  such that Anna's hearers reasonably believe that Tim is at least  $n$  mm tall, it's not a fact we're in a position to know. So there is no way for the participants in this conversation to coordinate on a single proposition in the cloud with which to update the common

ground after Anna's claim (25) has been accepted.

In the general case, then, we cannot expect that an exchange between speaker and hearer will lead to convergence on an update to a standard, Stalnakerian common ground (a set of worlds). That this can happen in the case epistemic *might* is due to special features of that case, which don't even carry over to *probably*.

## 7 TOWARDS A BETTER SOLUTION

The cloudy picture is an attempt to understand how communication can work in the presence of underdetermination. It is conservative in allowing us to hold on to two orthodoxies: the view that propositions have truth values relative to possible worlds, and the view that the common ground can be modeled as a set of such propositions (or alternatively a set of worlds), and perhaps a nonfactual "scoreboard." But being conservative in these respects requires a radical new picture of the speech acts involved in communication. On the orthodox picture, an assertion can be viewed as a proposal to add a single proposition to the common ground, and accepting an assertion is accepting this proposal. As we have seen, the cloudy view needs to replace this picture with an unorthodox new picture. We have looked at three different proposals for what this new picture could look like, and found all of them inadequate.

A better approach is to remain conservative in our conception of speech acts, and innovate in our conception of propositions. A first step is to allow propositions to vary in truth not just with a possible state of the world, but with one or more nonfactual parameters. In the case of vague gradable adjectives, this might be a delineation function that provides cutoff points; in the case of epistemic modals, it might be an information state. If we do this, we can keep the idea that the content of an assertion is a single proposition, and that the proposed update is to add this proposition to the common ground. This immediately solves the major problem facing Buchanan's and von Stechow and Gillies' views: that the hearer is not targeting a single possible-worlds proposition any more than the speaker is. For we now have a single content that is asserted by the speaker and recognized by the hearer.

We also solve the problem we saw with Braun and Sider's version of the cloud view. Recall that on their view, the standard for assertion is "approximate truth," or truth on every legitimate disambiguation, so if an assertion is accepted, the common ground must be updated with the conjunction of the legitimate disambiguations. That led to the in-

correct prediction that an assertion of (11), in a context where it is common ground that Richard is between 190–195 cm in height and the range of legitimate thresholds for *tall* is 185–195 cm, is to add to the common ground that Richard is 195 cm tall. Intuitively, though, the update does not rule out any particular height for Richard: it just rules out certain *combinations* of thresholds for *tall* and heights for Richard—those at which Richard’s height falls below the threshold. To accept an assertion of (11) is to accept a constraint on *combinations* of heights and delineations: to commit ourselves to Richard’s height being above the threshold for *tall*. In doing this, one may not rule out any particular heights or delineations, but only certain *combinations* of heights and delineations. So we need to think of the common ground as a constraint on *combinations* of worldly states of affairs and delineations—a set of (world, delineation) pairs. And once we do that, it is natural to model the content of (11) the same way. Update, then, becomes simple intersection.

There remains a problem of interpreting this formalism. We know what it is to accept an orthodox proposition—one that is true or false at a possible world—but what is it to accept one of these souped-up contents? And what is it for a set of (world, delineation) pairs to be common ground in a conversation? Until these questions are answered, we do not know how to apply our formalism.

Barker (2002), who should be credited with the basic formal insight that updates constrain combinations of delineations and worldly states of affairs, proposes that when we rule out delineations, we are reducing our “ignorance” about “the” delineation governing our discourse. Indeed, since a single delineation is determined by the world, we can take it to be a feature of the world and just use worlds instead of (world, delineation) pairs. We are back to an orthodox conception of propositions.

But this, I think, is a mistake. When there are multiple delineations in play, that is not because we’re ignorant of some fact of the matter, but because we haven’t made up our minds about a practical matter—where to draw the line for *tall*. This is reflected in the fact that, in a borderline case, it sounds funny to say

(27) Tim is probably tall,

but fine to say

(28) We should probably count Tim as tall.

We should think of contents and the common ground, then, as a set of pairs, with one element representing the content of a maximally determinate belief, and the other repre-

senting the content of a maximally determinate plan or intention. The idea should feel familiar: it is, more or less, the conception of contents in in Allan Gibbard’s plan expressivism (Gibbard 2003). These contents represent mental states that combine doxastic and practical elements, and that have, in general, a dual direction of fit—partly world-to-mind, partly mind-to-world.

Gibbard thinks of the second component of his pairs as *hyperplans*: fully determinate contingency plans covering every possible circumstance, and resolving all indecision. Since a hyperplan includes a resolution to put the threshold for *tall* in a particular place, we can think of our contents as (world, hyperplan) pairs instead of (world, delineation) pairs. (This also saves us from the possibly hopeless task of introducing separate parameters for information states, group boundaries, completions of definite descriptions, and everything else that is subject to felicitous underdetermination.) The problem of interpreting our formalism, then, reduces to the existing problem of interpreting the doxastic-practical states posited by plan expressivists.<sup>9</sup>

Appeals to clouds of proposition are motivated by the desire to retain an orthodox conception of propositions and an orthodox conception of the common ground. But this means finding an alternative to the orthodox story about assertion and communication. We have looked at three such alternatives and found them wanting. Better, then, to keep the orthodox story about assertion and communication, and adopt an unorthodox conception of propositions and an unorthodox conception of the common ground: roughly, the one advocated by plan expressivism.<sup>10</sup>

---

<sup>9</sup>Actually, in some respects the problem is easier. Gibbard cannot think of his hyperplans as fully determinate intentions, which would map circumstances onto particular actions, or he faces the “negation problem” (Gibbard 2003, 54–56; Dreier 2009; Schroeder 2008). To avoid this, he thinks of hyperplans as mapping circumstances onto sets of actions—intuitively, the actions that are permissible (Gibbard 2003). But this is problematic, since Gibbard’s aims require that he identify the mental states in non-normative terms (see Dandelet 2018 for an insightful discussion). For our purposes, though, we *can* think of the hyperplans as simply fully determinate intentions, which map circumstances onto particular actions.

<sup>10</sup>I am grateful for support from a fellowship at the Paris Institute for Advanced Studies, with the financial support of the French State managed by the Agence Nationale de la Recherche, programme “Investissements d’avenir,” (ANR-11-LABX-0027-01 Labex RFIEA+), and the Fondation Maison des Sciences de l’Homme. Thanks also to my Spring 2018 graduate seminar at Berkeley, and colloquium audiences at Chicago, Konstanz, Paris, Berkeley, Syracuse, and Stanford.

## REFERENCES

- Barker, Chris. 2002. "The Dynamics of Vagueness." *Linguistics and Philosophy* 25: 1–36.
- . 2009. "Clarity and the Grammar of Skepticism." *Mind and Language* 24: 253–73.
- . 2013. *Inquiry* 56: 240–57. <https://doi.org/10.1080/0020174X.2013.784482>.
- Braun, David, and Theodore Sider. 2007. "Vague, so Untrue." *Nous* 41: 133–56.
- Buchanan, Ray. 2010. "A Puzzle About Meaning and Communication." *Nous* 44: 340–71.
- Cartwright, Richard. 1962. "Propositions." In *Analytic Philosophy*, edited by R. Butler. Vol. 1. Blackwell.
- Dandelet, Sophie. 2018. "Preference and the New Negation Problem."
- Dreier, James. 2009. "Relativism (and Expressivism) and the Problem of Disagreement." *Philosophical Perspectives* 23: 79–110.
- Egan, Andy, John Hawthorne, and Brian Weatherson. 2005. "Epistemic Modals in Context." In *Contextualism in Philosophy*, edited by Gerhard Preyer and Georg Peter, 131–68. Oxford: Oxford University Press.
- Fintel, Kai von, and Anthony S. Gillies. 2011. "'Might' Made Right." In *Epistemic Modals*, edited by Andy Egan and Brian Weatherson. Oxford: Oxford University Press.
- Gibbard, Allan. 2003. *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Kennedy, Christopher. 2007. "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives." *Linguistics and Philosophy* 30: 1–45.
- Khoo, Justin, and Joshua Knobe. 2018. "Moral Disagreement and Moral Semantics." *Nous* 52: 109–43. <https://doi.org/10.1111/nous.12151>.
- King, Jeffrey C. 2014. "The Metasemantics of Contextual Sensitivity." In *Metasemantics*, edited by Alexei Burgess and Brett Sherman, 97–118. Oxford: Oxford University Press.



- . n.d. “Strong Contextual Felicity and Felicitous Underspecification.” *Philosophy and Phenomenological Research*.
- Lewis, David. 1979. “Scorekeeping in a Language Game.” *Journal of Philosophical Logic* 8: 339–59.
- MacFarlane, John. 2003. “Epistemic Modalities and Relative Truth.” <https://johnmacfarlane.net/epistmod-2003.pdf>.
- . 2011. “Epistemic Modals Are Assessment Sensitive.” In *Epistemic Modals*, edited by Andy Egan and Brian Weatherson. Oxford: Oxford University Press.
- . 2016. *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford: Oxford University Press.
- Price, Huw. 1983. “Does ‘Probably’ Modify Sense?” *Australasian Journal of Philosophy* 61: 396–408.
- Schroeder, Mark. 2008. *Being for: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford University Press.
- Stalnaker, Robert. 2002. “Common Ground.” *Linguistics and Philosophy* 25: 701–21.
- Suikkanen, Jussi. n.d. “Contextualism, Moral Disagreement, and Proposition Clouds.” *Oxford Studies in Metaethics*. <https://philarchive.org/archive/SUICMD>.
- Yalcin, Seth. 2007. “Epistemic Modals.” *Mind* 116: 983–1026.