

# Stupefying\*

## Abstract

People often accept assertions which they do not understand, a phenomenon I call *stupefying*. I argue that stupefying is a means for conversational manipulation that works through at-issue content, unlike the not-at-issue content and ‘backdoor speech act’ routes identified by Stanley (2015) and Langton (2018a). This shows that we need to reconsider a widely assumed connection between attention and at-issue content. It does not show, however, that we should avoid stupefying altogether, since stupefying turns out to have important cooperative uses, in addition to its manipulative ones.

Eben Byers—a Yale alum, steel magnate, amateur golf champion, and “ladies’ man”—suffered a minor injury from a fall on his train ride home from the 1928 Harvard-Yale football game. To assist his recovery and return himself to his pre-injury vigor, Byers followed the recommendation of his physician and started taking Radithor, a tonic of radium salts dissolved in water. Like many in the early 20<sup>th</sup> Century, Byers and his physician may have heard assertions of sentences like (1) made by promoters of radioactive cure-alls like Radithor.

- (1) The ionizing process of the Alpha rays sets up revitalizing forces in these glands [of aging], pouring renewed streams of hormones into the blood and bringing about the most astonishing results.<sup>1</sup>

Lacking anything but the vaguest notions of ionization, alpha rays, ‘revitalizing forces’, radioactivity, or what it means for a hormone to be ‘renewed’,

---

\*Acknowledgements removed for review.

<sup>1</sup>This comes from Morris (1926, p. 183), a book devoted to promoting radioactive medicines and Radithor in particular. It was ghost-written by William Bailey, owner and inventor of Radithor. Morris, a doctor, contributed a few pages and his name for the cover. For more of this story, see C. Clark (1997, Ch. 7).

these men would likely have just gone along with such statements, accepting without understanding them, especially since those making the assertions appeared to be well-informed doctors or scientists.

Sometimes accepting claims without understanding them doesn't have much effect beyond the conversation itself. But often it does. Having accepted (1), Byers and others would have been more prepared to accept claims like (2).

- (2) Ingesting radioactive radium water is beneficial to health and harmless in any quantity.<sup>2</sup>

Convinced, Byers took “some 1,400 bottles” of Radithor over the course of two years. In hindsight, it's not surprising that things did not end well. As a later *Wall Street Journal* headline put it, “The Radium Water Worked Fine Until His Jaw Came Off” (Winslow 1990).

*Stupefying* is what I will call this phenomenon of assertions being accepted without being understood. When a speaker asserts something an addressee doesn't understand, the addressee is put in a kind of stupor—an insensibility to the relevant implications of and evidence for or against accepting some claim—and partly because of this, for reasons to be explored, easily falls into taking on the claim as one of those jointly accepted in conversation. This article is a study of stupefying. We'll see that stupefying is a conversational tool with important uses, both helpful and harmful. We'll also see that the way it works is not easily captured by current theories of conversation.

Traditionally, philosophers of language have focused on conversations that are highly idealized, both descriptively and normatively. But there has been a recent surge of interest in conversations which are normatively non-ideal.<sup>3</sup> Besides allowing us to coordinate in intricate, flexible, and mutually beneficial ways, language also gives us means to manipulate others into acting against their own interests, doing things they would never have thought to do had we not conversed with them. How does it

---

<sup>2</sup>This is how the *Chicago Daily Tribune* (1932) reported it as being advertised.

<sup>3</sup>For some very recent work, see Langton (2018a) and Langton (2018b), Camp (2018), McGowan (2018), Saul (2018), Asher and Paul (2018), Tirrell (2018), Stokke (2018), D. Sosa (2018), Cappelen and Dever (2019), McGowan (2019), Beaver and Stanley (2019), Beaver and Stanley (forthcoming), Langton (forthcoming), Henderson and McCready (forthcoming), and Khoo and Sterken (forthcoming).

allow us to do this?

One important means of manipulation through conversation involves getting one's interlocutors to accept contents without first giving them due consideration. The work of Jason Stanley (2015) and Rae Langton (2018a) give us tools for thinking about this kind of manipulation. They point to various sneaky devices that operate in the conversational background, which can get content into the common ground without the other interlocutors attending to that content. Their picture, I believe, is a useful one, but it also suggests that this kind of manipulation can only happen through reliance on not-at-issue content or what Langton calls 'backdoor speech acts'. I'll argue that stupefying shows this cannot be right.

In §1, I present a preliminary account of the relevant kind of manipulation, bolstering Stanley's and Langton's ideas by bringing them together with recent work on the dynamics of attention in conversation. In §2, I argue that stupefying refutes this account, since it allows one to sneak explicitly asserted, *at-issue* content past one's interlocutor into the common ground—straight through the front door—without that content's being attended to. I then diagnose where the original theory went wrong and show how to revise it in light of the lesson from stupefying. In §3, I take up the question of why stupefying occurs and what should be done about it. It emerges there that while stupefying is often harmful, it can also be used harmlessly as part of fully cooperative and rational interactions, and indeed may be crucial for an important kind of instruction and learning. I conclude that while stupefying should be treated with caution, we should not adopt a blanket policy of avoiding it.

## 1 Short-Circuiting and Inattention

Conversations are had for various reasons and can accomplish various things. But if an exchange is to be a conversation, as opposed to mere talking-at, then the interlocutors must share the proximate goal of coordinating on a body of information which is to count as commonly accepted, or at least the goal of appearing to so coordinate.<sup>4</sup> This is why every serious model of conversation involves as a central component some

---

<sup>4</sup>For discussion, see Asher and Lascarides (2013).

notion of *common ground*: a body of information publicly accepted or committed to by the interlocutors at a given stage of the conversation.<sup>5</sup>

The cooperation required for conversation does not mandate the full alignment of aims, however. And indeed many conversations are characterized by competition, domination, and subversion. Often the only way to get others to act in some way, particularly when its at their expense, is through conversing with them. Speech, as Adam Smith put it, “is the great instrument of ambition, of real superiority, of leading and directing the judgments and conduct of other people” (Smith 1790, §VII.iv.25).<sup>6</sup> Manipulating the common ground to one’s favor is an important way of directing the judgments and conduct of others, an important way to control them. But how does this happen, exactly?

Sometimes, the most efficient way to persuade and direct others is to simply present your evidence and reasoning carefully and to be open to resolving various disagreements through rational means. When this is so, even a speaker with less than fully cooperative aims will behave as if they were interested simply in the results of joint rational deliberation. But often this is not the most efficient way of using language to influence others. Perhaps you don’t have overwhelming evidence for what you wish to persuade others of. Or you do, but don’t trust your audience to recognize it. Or you think they could recognize it, but it would take too long for them to do so. For these and other reasons, people often turn to other means of persuasion and influence with speech, means contrary to ideals of mutual respect. Sometimes this will involve getting one’s interlocutor to allow something into the common ground which they have not deliberated about sufficiently. This is what I will call *conversational short-circuiting*.<sup>7</sup>

What methods of short-circuiting are there, and how do they work? In

---

<sup>5</sup>See §1.1 below for a more detailed account of what common ground is.

<sup>6</sup>See also the title of Newt Gingrich’s 1994 GOPAC memo, “Language: A Key Mechanism of Control”. Readers of Stanley (2015) will not be surprised to find that “welfare” is on the list of negative words the memo recommends associating with one’s political opponents.

<sup>7</sup>Short-circuiting is a way of manipulating with speech, but it’s not the only way. For example, one can manipulate through lying or misleading, which undermines joint deliberation in another way. There are also more brute force methods, which don’t have much to do with joint deliberation: one can goad with insults and slurs, intimidate with threats, and so on.

the remainder of this section, I will develop a theory of an important kind of short-circuiting, *inattentional short-circuiting*, which involves adding content to common ground without one’s interlocutors attending to that content. The resulting account is plausible, but in §2 I will argue that stupefying shows it needs to be revised.

In brief, the account is as follows. Utterances can add contents to common ground in two ways: through *at-issue content*, which is in the conversational foreground, and *not-at-issue content*, which is in the background. Content which is at-issue is publicly attended to before it is added to common ground, making it an unpromising channel for short-circuiting. Rational interlocutors will not let information they are attending to into the common ground unless it is in their interest to do so. Not-at-issue content, on the other hand, need not be attended to before being added to the common ground, allowing for inattentional short-circuiting. We would expect, then, short-circuiting to operate through not-at-issue updates to common ground, such as presupposition accommodation, rather than through at-issue content.

## 1.1 Common Ground and Conversational Tone

To explain in more detail how content is snuck into the common ground, we’ll need to say just what common ground is. We’ll work with the view, due primarily to Stalnaker, that the common ground is the set of propositions which all of the conversational participants accept in a suitably public way.<sup>8,9</sup>

---

<sup>8</sup>See Stalnaker (1970), Stalnaker (1974), Stalnaker (1978), Stalnaker (1998), Stalnaker (2002), and Stalnaker (2014).

There are reasons to take common ground to be not a set of propositions but rather a set of probability spaces (Yalcin (2012), Moss (2018, Ch. 2)). We may also want common ground to extend beyond the doxastic, including questions under discussion and some preference-like ordering (See Murray and Starr ([forthcoming](#)) and references therein). Our purposes will be just as well served by the simpler picture, however.

<sup>9</sup> There is an alternative picture of common ground, which I’ll call the public commitments model, which stems from Lewis (1979) (though see also Hamblin (1971)). Recent work in this tradition includes Ginzburg (2012), Lepore and Stone (2015), and Farkas and Roelofsen (2017). On this view, we can treat a conversation as having a conventionally determined ‘score’, which can be updated in various rule-governed ways. Among other components, the conversational scoreboard will include the propositions which the participants are all jointly committed to: the common ground (or as some call it, the conversational record). There are ways to fill out this picture of common ground

What it takes for acceptance be suitably public may that among the acceptance is part of the interlocutors' common knowledge or belief (in the sense of Lewis (1969) and Schiffer (1972)), or it may be something less involved.<sup>10</sup> And what does it take to accept a proposition? Following Stalnaker, we'll take this to be a rather thin notion, where one accepts a proposition iff one treats it as true for the purposes of the conversation. One can accept by believing, but acceptance is weaker than belief, in that it includes pretense, supposition, and so on. Not all conversations are serious, sober ones where everyone is taken to believe what they say, even though by saying it they're acting as if it is true for some purpose. In these cases, contents are accepted, though not believed, but can be common ground nevertheless.

This introduces a puzzle for us, though. We're interested in how making additions to common ground could be used to control others. But on its own, something's counting as common ground in a conversation doesn't seem to mean much. Acceptance doesn't require belief, so one may accept something, act as if it is true for the purpose of the conversation, then go on one's merry way believing the opposite.<sup>11</sup> If Bob confronts Anne with the fact that she was committed in some conversation to  $p$  but now seems to be acting otherwise, Anne can sometimes easily explain herself by saying "Oh, we were just messing around—I didn't actually believe any of that" or "We were just talking about what would be the case supposing that  $q$ . I don't know whether  $q$  is true, though, and haven't

---

which tie it closely to the Stalnakerian picture. The relevant convention for something counting as a joint commitment, for example, could just be that it is publicly accepted. But common ground on this account can also be detached from the Stalnakerian version, by taking there to be alternative rules for determining the joint commitment part of the score. My arguments do not require us to decide between these understandings of common ground. Readers who prefer one over the other may take their pick. For those prefer, as I do, to follow Camp (2018) in incorporating both kinds of common ground into our theories of discourse, we may treat my claims about 'common ground' hereafter as referring to a strengthened notion, including only those contents which are common ground in the sense of being publicly accepted *and* common ground in the sense of appearing as joint commitments on the conversational scoreboard.

<sup>10</sup>For arguments that we should make due with less than common knowledge, see H. H. Clark and Marshall (1981) and Lederman (2018).

<sup>11</sup>And however we spell out what a conversational commitment is on the public commitments model (see note 9 above), in order to allow for the cases of pretense, supposition, and so on, it cannot in general be something strongly binding, enforced through social censure outside of the conversation.

made up my mind about  $p$ ". How, then, does getting something into the common ground make for an effective means of influence?

The answer: not all conversations are ones of pretense, or reasonably regarded as such. Anne's above excuses won't get her out of what she said or agreed to in any ordinary conversation, and it would be rather a poor defense of giving perjurious testimony under oath. We can draw here on what Yalcin (2007, p. 1008) calls *conversational tone*: the attitude which it is commonly known by the interlocutors to be the attitude which they will all take towards the content in the common ground.<sup>12</sup> In many conversations, it is public information that everyone will come out of the conversation believing what has been established there, unless new evidence appears. So if one gets something into the common ground where the conversational tone is belief, one will have succeeded in "directing the judgments and conduct of other people".

## 1.2 At-issue/Not-at-issue Updating

Getting something into common ground, then, can have real influence. But how does one do it? The paradigm way content is added to common ground is through assertion by a speaker and assent by the other interlocutors (or in some cases, arguably, absence of dissent).<sup>13</sup> There may be various means of making assertions, but the usual way it happens in spoken English is by a speaker uttering a declarative sentence with falling intonation. And assent from the addressee is indicated by an affirmative answer particle (*yes, okay, right*, and so on) or gesture (like nodding), or simply by continuing the conversation in a way that doesn't reject or

---

<sup>12</sup>This definition might require some tweaking. As with the propositions of common ground, requiring tone to be commonly known may well be too strong, in which case we can replace it with a weaker publicity requirement. And public *to the interlocutors* may be too weak, since it may be important that others not directly involved in the conversation are aware of the conversation's status. Finally, I think we should not require a uniform tone for the conversation, but instead should take certain parts of the common ground to be tagged with one tone, and other parts with others. We may also want the level of commitment to a particular proposition vary across interlocutors.

<sup>13</sup>Indeed, according to many, this effect on common ground is an essential effect of assertion: an assertion *just is* a proposal to update the common ground by adding content to it. For discussion, see Stalnaker (1999, p. 10), Farkas and Bruce (2010), MacFarlane (2011), and Stalnaker (2014, Ch. 2.5). We need not take a stand on whether this makes for a good analysis of assertion. It suffices to allow that this is the typical actual effect of an assertion when followed by acceptance by the addressees.

otherwise resist the assertion.<sup>14</sup> If  $\phi$  is the sentence being asserted, what is proposed to be added to common ground is  $\llbracket\phi\rrbracket$ , the semantic content conventionally associated with  $\phi$ .<sup>15</sup> If the addressee accepts, into the common ground  $\llbracket\phi\rrbracket$  goes.<sup>16</sup>

But this picture is too simple. Often there are multiple contents semantically encoded in a single sentence which have importantly different statuses. Compare the following:

- (3) a. The modern antivaccine movement can be traced to a paper published 16 years ago in *The Lancet*, a respected medical journal.<sup>17</sup>
- b. *The Lancet*, which published a paper 16 years ago to which the modern antivaccine movement can be traced, is a respected medical journal.

Both sentences encode the proposition that *The Lancet* is a respected medical journal and the proposition that the modern antivaccine movement can be traced to a paper published in the *Lancet* 16 years prior to the utterance. But they differ in how they present these propositions. In (3-a), the information about the origins of the antivaccine movement is presented as the main issue, whereas status of *The Lancet* is backgrounded. This can be seen by considering how (3-a) sounds in response to a question like “Where does the antivaccine movement come from?” (good) and how it sounds in response to a question like “What is *The Lancet*?” (bad). Things are just the reverse with (3-b). Potts (2005) introduces the term *at-issue*

---

<sup>14</sup>See Farkas and Bruce (2010). For alternative ways of resisting that don’t involve rejection, see Bledin and Rawlins (2016) and Bledin and Rawlins (forthcoming).

<sup>15</sup>I’m ignoring context-sensitivity here. More accurately, but still overly simplified: the value of the semantic character conventionally associated with  $\phi$  given the context of utterance as an input, as in Kaplan (1989). We might also allow the assertoric content—what gets proposed for addition to the common ground—to come apart from the semantic value of the sentence, even relativized to context. Then what gets added to common ground is a content which is not identical to the semantic value of the sentence, but some content determined by it (perhaps only partially). See Dummett 1973, pp. 446–447; Lewis 1980; Stanley 1997; Ninan 2010; Rabern 2012; Stalnaker 2018 though see also King 2007, Ch. 6; Cappelen and Hawthorne 2009; Stojnić 2017 for pushback.

<sup>16</sup>Many follow Ginzburg (1996), Roberts (1996/2012), and others in including in the model a component for recording what has recently been proposed for common ground-hood, as a kind of staging ground. Farkas and Bruce (2010) call this the Table.

<sup>17</sup>From the COCA corpus: <https://corpus.byu.edu/COCA>.



*content* for the information foregrounded in this way and *not-at-issue content* for the backgrounded information. In light of this, we need to make a revision to our view of assertion: it is the at-issue content of an assertion which is explicitly proposed to be added to common ground.<sup>18</sup>

I'll take not-at-issue content to include any other content conveyed by an utterance. This will be a quite broad and disunified class, including the contents expressed by supplements like nominal appositives and non-restrictive relative clauses (as in (3)), conversational implicatures, and presupposition triggers. Such breadth makes our next claim relatively easy to defend: some not-at-issue content gets added to common ground.

There are various controversies over just what not-at-issue content gets added to common ground and how it does so, but these will not concern us. All we need is that it sometimes does happen. One relatively uncontroversial case of not-at-issue update to common ground is through presupposition accommodation.<sup>19</sup> When one utters (5), one presupposes that Mary smoked at some point.

(5) Mary stopped smoking.

That is, one treats it as already established—already part of the common ground—that Mary smoked. But one can use such an utterance in cases where it isn't already established. In such cases, the presupposition typically becomes common ground, perhaps with the pretense that it was

---

<sup>18</sup>For Potts, at-issue content must be semantically encoded. But others, following Roberts et al. (2009) have drawn the distinction solely in terms of whether the content can be used to address the question which is the current focus of the conversation. This allows for at-issue content to include certain conversational implicatures, as in:

- (4) A: Are you coming to the party tonight?  
B: I've got a lot of work to do and have to get up early tomorrow.

Here the question which *B* is addressing is whether she is coming to the party, but the negative answer is not semantically encoded, but rather inferred pragmatically. On the Roberts et al. account, this is still at-issue content. The literal, 'what is said' content of *B*'s utterance will be not-at-issue. Similarly for cases of sarcasm, metaphor, and so on.

To make my claim about stupefying as interesting as possible, we can use a strengthened, hybrid notion of at-issueness. A content will count as at-issue for us when it is at-issue in Potts's sense *and* at-issue in the Roberts et al. sense. This rules out, on the one hand, the conversationally implicated contents which *do* address the relevant question, and semantically encoded contents which don't, on the other, making at-issue relatively narrow and not-at-issue quite broad.

<sup>19</sup>See Karttunen (1974), Stalnaker (1974), Lewis (1979), and von Stechow (2008).

there all along.

Interestingly, accommodation seems to happen by default. Presuppositions are not always accommodated; one can challenge them or outright reject them. But without any explicit challenging, they will typically go into the common ground, without any effort from the participants of the conversation, and regardless of what happens with any at-issue content of the utterance.

Even in cases where the at-issue content is rejected, the presuppositions will often be accommodated anyways.

(6) A: Has Mary stopped smoking?

B: I don't know, I didn't even realize that she did smoke.

Here the issue of whether Mary has stopped smoking is left unresolved, but the issue of whether Mary ever smoked is tacitly settled, as von Stechow puts it, "quietly and without fuss".

The distinction between updating the common ground through at-issue content and doing so through not-at-issue content (as with presuppositions) has been a crucial one to work on non-ideal speech. Stanley, for example, argues "that propaganda typically affects the not-at-issue content of an utterance. It enters into the common ground by routes other than assertion" (Stanley 2015, p. 172).<sup>20</sup> Langton argues that backdoor speech acts—"low profile speech acts, enabled by presuppositions and their ilk, that tend to win by default" (Langton 2018a, p. 146)—are an effective tool for 'evil' speech, since they can be difficult to block. This is in part because "[b]ackdoor speech acts can have an under-the-radar quality absent in assertion. When someone says 'Even George could win, attention is on 'George could win', not on 'even', still less on 'George is an unpromising candidate'" (Langton 2018a, p. 159).<sup>21</sup>

The idea, then, is that we can sneak content into the common ground through the not-at-issue backdoor. Asserted, at-issue content, however, lacks this ability. In the remainder of the section I'll flesh out this idea more explicitly.

---

<sup>20</sup>Stanley is using not-at-issue content in a somewhat narrower sense than I am. This will not affect the account I am spelling out here, however.

<sup>21</sup>For related proposals, see Haslanger (2011), McGowan (2004), McGowan (2019), and Saul (2018).

### 1.3 The At-issue/Attention Link

The account turns on what content speakers are aware of. To track this, we can enrich standard models of discourse with a class of agent-relative modal operators  $A_i$ , where  $A_i\phi$  says that the agent  $i$  is aware of  $\phi$ .<sup>22</sup> But when does an agent count as ‘aware’ of a sentence? As we’ll see in §2, there are a few related notions of awareness that will be useful to include in our theories, but for now we’ll follow Franke and de Jager (2011) and others in taking the relevant kind of awareness to be attention. More specifically, we’ll think of it as attention to the content of the sentence; attending to the mere phonetic or orthographic object won’t suffice.  $A_i\phi$ , then, says that  $A$  is attending to  $\phi$ ’s content, where what it takes to attend to a content is to be actively thinking or entertaining it. I attend to the content of “It’s raining” by attentively considering whether its raining or entertaining the thought that it is raining.<sup>23</sup> Even more specifically than that, we want  $A_i\phi$  to indicate attention to the content of the  $\phi$  as the content of  $\phi$ . If  $i$  is thinking some content, this won’t suffice for  $A_i\phi$  if that content just happens to be  $\phi$ ’s content.<sup>24</sup> I’m not attending to the content of a sentence, in the relevant sense, if it’s written in a language I don’t understand but by coincidence entertaining the content it actually expresses. Nor am I attending to its content in the relevant way if I am considering the question whether this sentence, whatever it means, is true. For  $A_i\phi$  to hold,  $i$  must be entertaining the content of the sentence itself, as the content of that sentence.

This is close to what we need, but it’s still not quite right. As we’ve

---

<sup>22</sup>The semantics of this operator can be given by adding to a discourse model functions  $\mathcal{A}_i$  from worlds to sets of sentences, specifying which sentences the agent  $i$  is aware of at a given world. See Fagin and Halpern (1988). In order to allow for awareness to change during a conversation, we should also index these operators to times, but I will ignore that here. There are various other, more informative ways we might wish to model agents’ awareness states (for example, with a question-based model like that of Yalcin (2016) and Bledin and Rawlins (forthcoming)), but we don’t need the extra structure for current purposes.

<sup>23</sup>What it takes for an activity—be it considering or something else—to be done attentively is an important topic in cognitive science and philosophy of mind, one which we cannot pursue here. See Mole (2011), Mole (2017), and Watzl (2017) for some options.

<sup>24</sup>We might try spelling this out as attending to the content because the sentence was uttered, but this will be subject to problems of “deviant causal chains”. Whether it can be spelled out some other way is an interesting question, but for now we’ll not try to define it in other terms.

seen, sentences often have multiple contents, so it can't be attention to 'the' content of the sentence that we're after. What we want instead is something which says which content, if any, an agent is attending to. One way to do this is by indexing our attention operators to contents as well as individuals, like so:

$$A_i^c\phi \quad \text{iff} \quad \begin{array}{l} \text{(i) } c \text{ is a content of } \phi \text{ and} \\ \text{(ii) } i \text{ is attending to } c \text{ as a content of } \phi. \end{array}$$

So for each content  $c$  that  $\phi$  has (and for each individual  $i$ ), there will be an attention operator for that content (and individual), that says the agent is attending to that content of the sentence.<sup>25</sup>  $A_i^{[\phi]}\phi$ , for instance, says that  $i$  is attending to  $\phi$ 's at-issue content,  $[\phi]$ . And if  $\pi$  is a presupposition of  $\phi$ , then  $A_i^\pi\phi$  says that  $i$  is attending to that presupposition.

Now that we've singled out the relevant notion of attention, we need to link it to what happens in a conversation.<sup>26</sup> As various authors have observed, it's hopeless to look for an account which fully specifies how attention shifts in response to an utterance. What people attend to varies dramatically from person to person and context to context, and not in any reasonably systematic way. However, there's one claim that seems to be generally agreed upon by those who have worked on the dynamics of attention in conversation: asserted content is publicly attended to.

Bledin and Rawlins (2016), for example, rely on a principle called *Drawing Attention*, which says that if a speaker utters  $\phi$ , then the attention state of every other discourse participant  $a$  is immediately updated to ensure  $A_a^{[\phi]}\phi$ , prior to any further processing (such as allowing into the common ground or challenging the assertion in some way).<sup>27</sup> For

<sup>25</sup>And we make the corresponding adjustment to our discourse models, now with functions  $\mathcal{A}_i^c$ , still from worlds to sets of sentences, specifying which sentences the agent  $i$  is aware of having  $c$  as a content at a given world.

<sup>26</sup>There are various principles we might want to use to constrain  $A_i^c$ , given that we wish to interpret it in this way. It seems natural to suppose that when an agent attends to the at-issue content of  $\phi$ , they also attend to that of  $\neg\phi$  (and vice versa). Whenever you're considering whether  $p$  is true, you're also considering whether  $\neg p$  is true. We might want to require that when an agent attends to the at-issue content of  $\phi \wedge \psi$ , they also attend to that of  $\phi$  and of  $\psi$ . For several authors working with the notion of attention, it is important that if one attends to the content of  $\diamond\phi$ , one also attends to that of  $\phi$ . We need not settle here what principles we should impose as axioms for how  $A_i^c$  behaves. All we need is a principle connecting  $A_i^c$  to conversational moves.

<sup>27</sup>Modulo their notation and treatment of attention as having subject matters, rather than propositions, as its content.

Franke and de Jager (2011), it's an important that "unawareness from inattentiveness is lifted spontaneously whenever agents process linguistic information that contains mention of an unaware contingency". Others concur.<sup>28</sup> Accepting any such proposal will imply the following principle:

**At-issue/Attention Link:** At-issue content of an utterance is always attended to by each interlocutor before being added to common ground.

What about non-at-issue contents? Some, no doubt, will in fact be attended to. And perhaps certain ones, like those associated with appositives (as in (3)) will be attended to as automatically as at-issue content is.<sup>29</sup> But often they are not. Presuppositions, for instance, are not always attended to before they are added to the common ground. This was just Langton's point. There is thus no strong link between not-at-issue contents and attention.

**Not-at-issue/Attention Separation:** Not-at-issue content is *not* always attended to by each interlocutor before being added to common ground.

What results from these two principles is a plausible account of how conversational short-circuiting can happen. One way to ensure that content is not sufficiently deliberated about before it is added to the common ground is through inattentional short-circuiting—getting it into the common ground without your interlocutors attending to that content. If At-issue/Attention Link is true, then inattentional short-circuiting cannot be done with at-issue content. However, given Not-at-issue/Attention Separation, it *can* be done with not-at-issue content. Like sleight of

---

<sup>28</sup>Roelofsen (2013, p. 194) says that in uttering  $\phi$ , a speaker is taken to "Draw attention to all the possibilities in  $[\phi]$  as possibilities that may contain the actual world." Crone (2017) assumes that "an utterance minimally raises awareness of all atomic proposition letters within that utterance". Similar ideas appear in Ciardelli, Groenendijk, and Roelofsen (2009, p. 98), Crone (2018), and Bledin and Rawlins (forthcoming).

<sup>29</sup>That said, it's worth noting that with appositives, the timing of the attention might differ from that of at-issue content. And it seems that when it comes to be the addressee's turn to speak, the attention may be required to be only on the at-issue content. Subsentential attentional dynamics is yet to be explored, but it seems to me that it might account for the different discourse behavior of sentence medial and sentence final appositives observed by Syrett and Koev (2014) and Anderbois, Brasoveanu, and Henderson (2015).

hand, inattentional short-circuiting relies on misdirection of attention, the oldest trick in the book.<sup>30</sup> While you attend to one thing—like the at-issue content—the speaker has snuck something else into the common ground. Coming out of the conversation, at least if it had the right conversational tone, you may find yourself with new commitments you didn't consent to and new beliefs (if only implicit ones) to which you never gave consideration.

If this is right, then there's an important means of conversational short-circuiting which is only available for contents packaged in a not-at-issue way. We can see why not-at-issue content would be a more promising vector for short-circuiting, just as Stanley and Langton suggest. It is thus unsurprising that not-at-issue content has been the focus of much of the work on manipulative speech. This is an attractive account, but as we'll see in a moment, it is wrong.

## 2 Stupefying as At-Issue Short Circuiting

In this section I argue that stupefying forces us to give up the At-Issue/Attention Link. It is a means of inattentional short-circuiting that works with at-issue content.

### 2.1 Two Kinds of Stupefying

Stupefying involves an addressee accepting an assertion without understanding it. When Eben Byers accepted claims like (1) about why radium water would help him, he was stupefied. And not only did he not understand what was said, but he *couldn't* have understood what was being said. Having no idea what ionization or alpha rays are, he lacked the relevant concepts to grasp the sentence's content. This kind of stupefying, which arises from lacking the conceptual resources to grasp the asserted content is what I'll call *stupefying by content*.

Sometimes we are stupefied not because we can't grasp the relevant content, but because we don't recognize that the assertion that's just been

---

<sup>30</sup>Various means of manipulating people to act against their own interests rely on misdirection of attention. See Akerlof and Shiller (2015) for some examples from economics.

made is a way of expressing that content. Suppose on your visit to the Outer Banks in North Carolina, you ask an Ocracoke Island local whether it's a good idea to stay at a particular cottage. They tell it is not, then explain:

(7) Its pizer is all whopperjawed.

Now, you might ask what a pizer is and what it is for it to be whopperjawed. But you may just as well say "Oh, okay," accepting (7) without understanding it, then go on to ask about the other place you were considering.

In this latter case, you've been stupefied, but not because there was some content which you were incapable of understanding. In the speaker's dialect, *pizer* just means *porch* and *whopperjawed* means *crooked* or *misaligned*.<sup>31</sup> You can, of course, grasp the proposition that the cottage's porch is crooked. Unlike the Byers case, it's (7)'s form rather than its content that prevents you from understanding it. This is *stupefying by form*.

Linguistic ignorance is not the only way stupefying by form can occur. Someone may have put things in such a convoluted way that processing it would take more effort than one can or is willing to expend. Or they may have simply mumbled. Again, in cases like these, one might follow up with a request for clarification, but often one instead accepts without understanding.

These two kinds of stupefying—stupefying by content and stupefying by form—map neatly onto ways the awareness operators like  $A_i^c$  have been interpreted in the economics and computer science literatures.<sup>32</sup> Above we noted that besides the attentional sense discussed in linguistics, it is sometimes understood in other ways. Indeed, usually the notion is introduced as being about "lack of conception" (Schipper 2015, p. 77), using examples like this one:<sup>33</sup>

How can someone say that he knows or doesn't know about  $p$  if  $p$  is a concept he is completely unaware of? One can imagine

---

<sup>31</sup>See Wolfram and Reaser (2014, Ch. 5).

<sup>32</sup>Beginning with Fagin and Halpern (1988). For a useful overview, see Schipper (2015).

<sup>33</sup>Though see Fritz and Lederman (2015, pp. 9–12). They argue that it's also the attentional sense that is typically the requisite one for applications in economics.

the puzzled frown on a Bantu tribesman's face when asked if  
he knows that personal computer prices are going down!

Fagin and Halpern (1988, p. 40)

This is very close to stupefying by content: one can just as well imagine the Bantu tribesman saying that yes, of course he knows that, despite not being able to grasp the relevant content due to lack of the relevant concepts.

The other way 'awareness' is sometimes interpreted is as an ability to process, on which "an agent is aware of a formula if he can compute whether or not it is true in a given situation within a certain time or space bound" (Fagin and Halpern 1988, p. 41). Work on this typically focuses on issues of computational complexity and resource bounded algorithms for computing truth in interpreted formal languages, but for application to natural language discourse one could incorporate linguistic ignorance, facts about human sentence processing, and features of the form of an utterance which can affect its interpretability in a given context. In this sense of awareness, you lacked awareness of (7), though you do have the concepts required to understand it, so are aware of it in the conceivability sense.<sup>34</sup>

Observe that both of these new kinds of unawareness imply the inattentional sense of unawareness we were working with in §1. One cannot attend to the content of a sentence in the relevant sense if one cannot process the sentence or if one cannot grasp the sentence's content. This means that when someone has been stupefied by either form or content, one accepts a sentence without attending to its content.<sup>35</sup>

---

<sup>34</sup>To keep track of these distinct notions of awareness, besides our attentional awareness operator  $A_i^c$ , we can introduce a similar operator for awareness in the ability to conceive sense,  $C_i^c$ , and awareness in the ability to process sense,  $P_i^c$ . What I've called stupefying by content, then, involves the addressee,  $a$ , publicly accepting  $\phi$  where  $\neg C_a^{[\phi]}\phi$ . And stupefying by form involves  $a$  publicly accepting  $\phi$  where  $\neg P_a^{[\phi]}\phi$ . If we take  $\neg C_i^c\phi$  to imply  $\neg P_i^c\phi$ , as I think we probably should, it may be more useful to define stupefying by form as stupefying *merely* by form, which would involve acceptance of  $\phi$  when  $\neg P_a^{[\phi]}\phi \wedge C_a^{[\phi]}\phi$ .

<sup>35</sup>We might take there to be a third kind of stupefying, which involves mere lack of attention. When one's mind wanders as someone talks, one may *be able* to process the relevant utterances and grasp their contents, but not actually do so because of distraction. In such conditions one sometimes accepts things absent-mindedly, without understanding what has been said. This could be defined as accepting  $\phi$  while  $\neg A_a^{[\phi]} \wedge P_a^{[\phi]}\phi$ .



And note that the unawareness of the stupefied concerns the at-issue content of those utterances which they are accepting. Byers' incomprehension stemmed from not having the concepts for what is expressed by "ionizing processes" or "alpha rays", the meaning of which contributes squarely to the at-issue content of (1). Because of this, he could not have attended to this content in the relevant sense. Because you didn't know the meaning of *pizer* and *whopperjawed*, you couldn't attend to the at-issue content of (7). And so, assuming that acceptance of these utterances resulted in additions of their at-issue contents to common ground, stupefying involves addition of at-issue contents to common ground without their being attended to. This would mean that the At-issue/Attention Link does not hold, and we should not expect inattentional short-circuiting to happen only through not-at-issue content.

## 2.2 Stupefying and Common Ground

We may doubt, though, that these contents really are added to common ground. After all—an objection could go—if you can't even grasp the relevant content how can you believe it? And if you don't believe it, how can it be in the common ground?

Recall that it's not belief that is necessary for common ground status, but acceptance. And while belief may require grasp of the relevant content,<sup>36</sup> acceptance should not.<sup>37</sup> We do need to have something that connects the agent to the content they accept, otherwise it would be a mystery why they would count as accepting this content rather than some other. But this can be satisfied in various ways, such as the addressee thinking of the content under the description "whatever the speaker just said, whatever it means". Similarly for whatever extra is required for making the acceptance suitably public. We can commonly believe that we all accept whatever the speaker just said, even if not all of us understand what it meant. So it is possible for there to be contents in the common ground that not everyone in the conversation can grasp.<sup>38</sup>

---

<sup>36</sup>Though see Recanati (1997).

<sup>37</sup>Stalnaker himself doesn't take acceptance of a proposition to require grasping it. In 1700, William III tacitly presupposes and so accepts that England can avoid nuclear war with France, despite not having the concept of nuclear war (Stalnaker 1984, pp. 88–89).

<sup>38</sup>This is also true on the public commitments model of common ground discussed in

Nevertheless, even granting that it's possible, we might still doubt that the at-issue content of (1) or (7) do in fact go into the common ground when addressees affirm them without understanding. Perhaps we can make do with less. Here are a couple options for doing so. We might take the *least common denominator view*, according to which all that goes into the common ground in these cases is whatever parts of the assertion's content that both the speaker and the addressee recognize (or perhaps: recognize in a suitably public way).<sup>39</sup> So in the case of (7), all you understand of the utterance is that the cottage has some thing with some property that makes it unsuitable for you to stay there. And all Byers gets from (1), presumably, is that it Radithor does something that will bring about astonishing results.<sup>40</sup>

Alternatively, we might adopt the *metalinguistic view*, which says that in cases of stupefying, what gets accepted isn't the at-issue content of the assertion, but rather an associated metalinguistic content, like (8).<sup>41</sup>

(8) That utterance of "Its pizer is all whopperjawed" is true, whatever it means.

With either the least common denominator or metalinguistic view, the at-issue content of the utterance doesn't go into the common ground. This would mean stupefying would not be a way of inattentional short-circuiting with at-issue content, since the content that the speaker can't attend to never makes it into the common ground. If this were right, we could retain the At-issue/Attention Link. We should reject these alternative views, however, since there is good reason to think that the at-issue content *does* get into the common ground.

There are various ways to act as if some content is true for the purposes

---

note 9 above—one can easily be committed to something one doesn't understand.

<sup>39</sup>What is a part of a content? We probably don't want it to be just any entailment, but rather one of the more demanding notions of content parthood from Yablo (2014) or Fine (2017).

<sup>40</sup>We might add to this view that often in cases of stupefying, the result is a defective context (Stalnaker 2002, pp. 717–718) where the speaker thinks there is more in the common ground than there really is.

<sup>41</sup>One way to implement this strategy is with Stalnaker diagonalization (Stalnaker 1978; Stalnaker 2004). Indeed, diagonalization is proposed by Stalnaker as a repair strategy for when someone says something that the other doesn't understand (though limited to cases where it is clear to all involved that the addressee wouldn't understand).

of a conversation. Some of these will involve understanding the content and acting in ways that would be rationalized by belief in such a content, like taking out an umbrella because one accepts that it will rain soon. But for conversation, the primary ways of acting as if some content is true are tied to linguistic conventions. One important way to act as if  $p$  is true is to assert something that means  $p$  (i.e., that has  $p$  as its at-issue content). Another is to assent to someone else's assertion of something that means  $p$ . In cases of stupefying, one assents to assertion of a sentence that means  $p$ , so one acts as if  $p$  is true for the purposes of the conversation. So one accepts that content. And this assent is suitably public, so one's acceptance is enough for to get that content into the common ground, given that the speaker also publicly accepts it. So we should reject views which imply stupefying doesn't get the at-issue content of the accepted assertions into the common ground.<sup>42</sup>

Here's another argument, primarily against the least common denominator view. Explaining presupposition licensing is one of common ground's central theoretical roles. One way to test for whether something has made it into common ground, then, is to check whether it can thereafter presupposed without accommodation. And stupefied acceptance licenses the presuppositions we'd expect if their at-issue contents had been added to the common ground. For example, (9) is a natural follow-up question

---

<sup>42</sup>Hawthorne and Magidor (2009, p. 394) make a similar point as part of their argument against Stalnaker's Uniformity constraint, which states that in cases of rational communication, an assertoric utterance expresses the same proposition in each possible world in the context set. Stalnaker (2009, pp. 408-409) and Almotahari and Glick (2010, p. 1085) worry that dropping this constraint will sever the connection between acceptance of a proposition in conversation from how the agent acts outside of the conversation. I do not share this worry, however. First, given what we've already said about conversational tone, it's clear that any connection between acceptance in conversation and action outside the conversation should be flexible and not too strong. Second, the connection between acceptance and action is not fully severed: in the *pizer* case, acceptance even by an uncomprehending addressee will have implications for how they will act both inside and outside conversation. In conversation, it affects what they can and can't presuppose. Outside conversation, it affects which cottage they will choose to rent. Finally, even if we follow Stalnaker in his ambitious Grice-inspired program to reduce linguistic communication to non-linguistic behavior and attitudes, we need not hold that *every* speaker in a linguistic community has behavioral dispositions sensitive to all the distinctions that the contents expressed in their conversations draw. For familiar content externalist reasons from Burge (1979) and others, we should try to get by with the much weaker requirement that the speech community as a whole is sensitive to the relevant distinctions.

you (or some other participant of the conversation) might ask after you've accepted (7).

- (9) Does the other cottage down the road have a whopperjawed pizer, too?

This presupposes that something salient besides the other cottage down the road has a whopperjawed pizer. Presumably it is licensed here because the proposition that the first cottage has one is already part of the common ground, rather than through presupposition accommodation. Having already accepted (7), it would be strange to challenge the additive presupposition of (9). So it is important that we take the common ground to have been updated with the at-issue content of (7), rather than something more general.

The objector would likely try to apply the same strategy to reinterpreting the effects of (9) that they applied to (7). The least common denominator proponent might claim that this question is just asking whether there's something wrong with the other cottage, and has a similarly broad presupposition, one that would be satisfied by what they take to have been added by (7). I think this cannot be right. Contrary to what this proposal predicts, (9) could not be correctly answered by (10).

- (10) Yes, its air conditioning is broken.

The metalinguistic view fares better here. The reinterpretation of (9) would presumably be that it is asking whether "has a whopperjawed pizer" applies to the other cottage, and presupposes that it applies to something else salient. (10) would be a bad answer to this question and with the acceptance of the metalinguistic interpretation of (7), the presupposition would be satisfied, as desired.

There's an additional reason to reject the metalinguistic view, however. According to this approach, when the addressee doesn't understand the at-issue content of an utterance, the at-issue content doesn't get added to the common ground, only metalinguistic information does. But understanding an utterance is a matter of degree. Consider Burge's famous case of a man who thinks he has developed arthritis in his thigh, but seems to have a number of other reasonable and correct beliefs about arthritis (Burge 1979, p. 77). When he was told by his doctor beforehand that he

has arthritis in his wrists, he no doubt understood what was said less well than his doctor did, but understood it better than what Byers could understand claims about alpha rays. We can imagine a range interlocutors with understanding of assertions declining gradually from the doctor's understanding of what they said, down to Burge's confused man, all the way down to the minimal understanding that Eben Byers and the visitor to Ocracoke Island had of the utterances they accepted, and there will be no clear point at which we go from an interlocutor who understands to one who doesn't.<sup>43</sup>

If we take the metalinguistic view we are thus left with two unappealing options. One is to say that the at-issue content is *never* what gets added to the common ground when an assertion is made and accepted, only the corresponding metalinguistic content is.<sup>44</sup> The other is to say that despite the smooth decrease, at some degree of understanding (perhaps not the same for all conversations), what acceptance does switches from adding the at-issue content to common ground to adding the associated metalinguistic content instead. Those taking this option would be stuck with the task of saying why there would be this sudden flip. Neither option is absolutely intolerable, but I take it to be a cost to have to accept either. Better to take stupefying, like other instances of assent to assertions, to involve addition of at-issue content to the common ground.

In my view, then, stupefying results in additions to common ground of contents, often at-issue contents, which have not been attended to by the addressee. So I conclude the At-Issue/Attention link is false, which is where the account of §1 went wrong. We can fix it by dropping it.

---

<sup>43</sup>This suggests a modification to the view of stupefying presented thus far. We should replace the binary awareness operators  $A_i^c$ ,  $P_i^c$ , and  $C_i^c$  with degree operators which return the degrees of attention, ability to process, and grasp of  $c$  that  $i$  has, respectively. As far as I know, gradable awareness operators have yet to be studied. Halpern and Piermont (2019), though, is a step in the right direction.

If we wanted to stick to our binary notion of stupefying, we could define it in terms of (contextually-variant) thresholds on these degree operators. Stupefying would occur when one accepts without understanding well enough for the purposes of the conversation. It might be better, though, to also take stupefying itself to be gradable—one can be more or less stupefied. Byers was highly stupefied, Burge's confused man less so, and a doctor who has been told that someone has arthritis even less.

<sup>44</sup>We can make this less unappealing by allowing other contents which the interlocutors (publicly) take to follow from the meta-linguistic content to get into the common ground as secondary effects. This makes better predictions, but is still intuitively unattractive. Conversation does not seem to be primarily in the formal mode.

This allows us to see stupefying as a potential means of inattentional short-circuiting which works for at-issue content. One can get at-issue content into the common ground without that content's being attended to, since you can get this content into the common ground without your interlocutor understanding it well enough to attend to it.

This seems to be what happened to Byers and his doctor—they didn't give due consideration to the claims promoting Radithor because they were stupefied by them. And if you decided not to stay at that cottage, you may have been stupefied by (7): perhaps a crooked porch isn't a drawback that would have deterred you from staying there, or perhaps you would have checked to see if it was crooked enough to be a problem, had you understood what was said. So it's not only not-at-issue content and backdoor speech acts we should watch out for. Content can be slipped into the common ground through the front door.

### 2.3 A Norm of Attention?

Though we should give up the At-issue/Attention Link, the idea behind it isn't totally off-track. It does seem that in accepting through saying *Okay* or *Oh, I see*, etc., one typically does act *as if* one understands and is attending to the relevant content, rather as one acts as if one believes or knows the content of what one asserts. How can we account for this without the At-issue/Attention Link?

One way is by appeal to a norm of correctness for acceptance, rather than some automatic updating of attention of the kind required for the At-issue/Attention Link. Lying shows us that lack of belief doesn't always prevent assertion, but we still think that there's some conversational norm which requires belief in or knowledge of what one asserts.<sup>45</sup> This is why, perhaps, asserting is a way of representing oneself as believing. Similarly, stupefying shows us that lack of attention doesn't always prevent acceptance, but we might still think that there's some conversational norm which requires attention to the content one accepts. This would explain why it seems that accepting is a way of representing oneself as attending. I suggest that we replace talk of automatic public attention with a normative

---

<sup>45</sup>Here I am thinking of norms of the kind discussed in Williamson (2000, Ch. 11) and the literature responding to it.

requirement.

Pursuing this line of thought, we could propose the following.

**Attentional Norm of Confirmation:**  $i$  may confirm  $\phi$  only if  $A_i^{[\phi]}\phi$ .

That is, confirm only if you are attending to the at-issue content of what you're confirming. We need not take this to be the only norm governing confirmation. You shouldn't confirm something you know to be false, for example. Nor do we need to assume that it is a fundamental norm; perhaps it can be derived from some other norm(s). However, it is the norm we'd need most directly for understanding why acceptance seems to represent one as attending.

Conversational participants, it seems, tend to try to follow this norm and often one is in some sense criticizable if one violates it. If it emerges that one had confirmed something one didn't understand, one can be chided: "Why did you agree if you didn't understand what I was saying?". However, I don't think that there can be such a norm in general. In some conversations, the following kind of hedged confirmations can be an acceptable response to assertions one doesn't understand.<sup>46</sup>

- (11) a. Hmm, I don't think I understand that, but okay.  
b. I have no idea what that means, but I'm sure you're right.

If the Attentional Norm held in general, we'd expect hedged confirmations like these to be always unacceptable, as the corresponding ones seem to be for putative norms of assertion.<sup>47</sup>

- (12) It's raining but I don't know/believe it is.

So while assertion may well have a norm of knowledge or belief, we should not accept that there is an attentional norm for confirmation.

It's important, though, that one can't say these things in just any circumstance. If you're told that it's best not to stay *at the house at with*

---

<sup>46</sup>Note that in conversations with hedged acceptances like this, it will be realized by everyone that stupefying has happened. This means that the common ground won't be defective, so we cannot treat stupefying in general as a special case of defective context.

<sup>47</sup>Interestingly, some ways of confirming don't sound okay when hedged like this. *Okay* and *sure* are suited for it, but *yes*, *of course*, *mhmm*, *definitely*, and others are not. So we should probably distinguish different varieties of confirmation. The Attentional Norm might hold for some of them.

*the whopperjawed pizer*, uncomprehending acceptance seems worse than accepting that you shouldn't stay at some given house which you can already identify, because it has a whopperjawed pizer. When something about the pizer is cited as a reason for some piece of advice, the uncomprehending addressee can still accept the advice, if only on faith. But when it is used to specify what the advice is, hedged confirmation doesn't work because one can't take the advice in the way the conversation seems to require.

This suggests that rather than appealing to a norm of acceptance, we should try to derive the normative requirement for attention in most cases from more general principles of pragmatics and rationality, in a roughly Gricean way. Something like:<sup>48</sup>

**Attentional Cooperative Principle:** Apportion your attention as is required by the accepted purpose or direction of the talk exchange in which you are engaged.

Like Grice's Cooperative Principle, this is very general and rather vague. Though it would take a lot of substantive work, the hope is that we could derive from it the requirement for most cases of attending to what one accepts, while still allowing for certain exceptions.<sup>49</sup> This would explain why accepting usually seems to involve acting as if one is attending—and why the At-Issue/Attention Link has been an attractive assumption to many—but nevertheless allows for the possibility of stupefied acceptance, even an openly stupefied acceptance.

Much work would be required to defend and flesh out the details of a proposal along these lines. To do so here would go beyond what space

---

<sup>48</sup>Based on the Cooperative Principle from Grice (1967, p. 26): "Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged". We could, of course, derive our principle from the Cooperative Principle if we make the somewhat unnatural move of taking one's attention to be a 'conversational contribution'. I take it, though, that Grice's principle is about how to be a good speaker. Our principle, by contrast, is primarily about how to be a good listener.

<sup>49</sup>More ambitiously, we might aim to account for the differences in felicity in inattentive acceptance when using different means of confirmation (see note 47). Still more ambitiously, we might try to extend this approach to account for the differences in felicity in using context-sensitive expressions (pronouns, possessives, quantifiers, tense, etc.) when the context underspecifies their semantic values, as discussed by Buchanan (2010), King (2018), and others.



allows, so I will leave that project for elsewhere. All I can say here is that it seems plausible that we can account for the connection between at-issue content and attention without committing to the At-Issue/Attention Link in a way that would explain why this Link would have looked plausible.

So far I have argued that stupefying is a way that at-issue content can be added to the common ground without being attended to, and so the At-Issue/Attention Link should be rejected. In the space remaining I will briefly explore some of the reasons why stupefying happens.

### 3 Why Does Stupefying Happen?

Why do speakers sometimes say what their addressees won't understand? And why do addressees sometimes accept it anyways? This final section is an exploration of some answers to these questions.<sup>50</sup> A full account of why stupefying happens is no more feasible than one of why lying happens. Nevertheless, I hope the following will clarify some of the major reasons stupefying happens and point us in interesting directions for further investigation.

#### 3.1 Why Speakers Stupefy

Stupefying is often done intentionally. Given that it can be a means of short-circuiting, it's no wonder that it is attractive to hucksters, quacks, frauds, and the like. It is a tool for manipulation.

One particular attraction of stupefying to the unscrupulous is that it sometimes allows for a kind of bootstrapping of epistemic authority. One method is what could be called *the Guru's Gambit*: confidently say some gibberish, using impressive sounding esoteric vocabulary. Imply this is connected to areas of deep concern to your audience—relationships, health, money, etc.—and that the gibberish gives you knowledge of feasible ways they can make big improvements in these areas. Since you're saying impressive sounding things which they think they can't understand, but

---

<sup>50</sup>An informal exploration. Some of these issues, I think, can be fruitfully studied in a game theoretic framework which incorporates unawareness (as in, e.g., Heifetz, Meier, and Schipper (2013), Halpern and Rêgo (2014), and Franke (2014)), and a public announcement logic (as in Ågotnes and van Ditmarsch (2011) and van Benthem (2011, Ch. 15)). I will leave this technical work for other occasions.

seem to make sense to you, you must know something they don't. You must be relatively expert on the matter of quantum healing or whatever. And as an expert, you may have some extra insight into what it can do. Moreover, since they don't understand what you're saying, no objections come to mind. One is inclined, then, to accept both the stupefying assertions and the more quotidian ones based on them, like "My method can help you" and "You should buy my book". This is of course helped along by a healthy dose of wishful thinking and, when others have already bought into it, the stupefied acquiescence of the surrounding crowd.<sup>51</sup>

A similar kind of bootstrapping can help explain obscurantism, which involves making statements which are difficult for one's audience to understand, but have contents which can be expressed in simple ways that could be easily understood. It is an attempt to stupefy by form while giving the addressee the impression that they are being stupefied by content. Like the Guru's gambit, obscurantism can trick an addressee into assigning the speaker more epistemic authority than they deserve. This is why charlatans often make needless use of jargon and put things in overly convoluted ways.

In cases like these, it rarely matters to the speaker what content, if any, the stupefying utterances have. All that matters is that it can successfully stupefy, then be leveraged in the right ways. This is why attempts to stupefy so often involve bullshit, in Frankfurt's sense of lacking a "connection to a concern with truth" (Frankfurt 1982, p. 125).

Intentional stupefying is not limited to the domain of posers and purveyors of bullshit, though. For various reasons, real experts sometimes deploy their knowledge in ways intended to stupefy their audience. Here is an example from Elena Ferrante's *Those Who Leave and Those Who Stay*. Lina is telling Lenù, the narrator, about her work with computers—newfangled technology at the time.

She talked to me about ferrite cores, rings traversed by an electrical cable whose tension determined the rotation, 0 or 1, and a ring was a bit, and the total of eight rings could represent a byte, and that is a character. . . . Is that clear? she asked me every so often. I answered yes, weakly, but I didn't know what

---

<sup>51</sup>See Sperber (2010) and Pennycook et al. (2015) for similar ideas. Langton (2018b) observes a related form of bootstrapping of authority in hate speech.

she was talking about. I perceived only that she noticed that nothing was clear to me, and I was ashamed of this.

Ferrante (2014, p. 262)

Lina is “bewildering [Lenù] with an incomprehensible jargon” (p. 261), apparently intentionally. It’s not obvious why she’s doing so, but one can imagine a range of motivations: to show off, to put down, to not insult by dumbing things down. Perhaps she’s just excited to show her friend that she’s finally found work that is enjoyable and intellectually engaging. She could well respond to Lenù’s shame by telling her it’s okay if she doesn’t understand and thanking her for listening.

Intentional stupefying can happen in various ways, for various reasons. But not all stupefying is intentional, and indeed it often goes unrecognized even after it has happened. Here’s a real example of unintentional stupefying of some concern to bioethicists. In his study of patient consent to medical research, John Fletcher describes the case of Mrs. B., a hospitalized patient who took part in a non-therapeutic study of dyslipoproteinemia. Fletcher reports:

... [Mrs. B.] complained in her interview with me that she lacked knowledge about the purpose of Dr. A.’s study, and that she had difficulty understanding the technical parts. When I reminded her that she did have ample opportunity in the consent process to ask such questions of Dr. A., she said “Well, I didn’t want him to think I was stupid!” She had pretended to understand some things and did not question Dr. A. as much as she wanted.

Fletcher (1976, p. 270)

Mrs. B. was, it seems, stupefied by some of the things Dr. A. said in their discussions leading up to her consenting to the study.<sup>52</sup> But this was not intentional on Dr. A.’s part. According to Fletcher, “Dr. A. actually

---

<sup>52</sup>Though this study didn’t involve high risks, it did prolong Mrs. B.’s hospital stay and involve injections as well as daily blood and urine collections. Had she properly understood the study and its purpose, perhaps she would not have consented to these inconveniences, or perhaps Dr. A. wouldn’t have wanted to spend enough time get Mrs. B. to understand the issues as well as she would have liked. Instead, Mrs. B. was apparently stupefied and deliberation was thereby short-circuited.


believed Mrs. B. to be very well informed and highly curious about the technical side of medicine" (Fletcher 1976, p. 272).

Experts often have a hard time assessing what non-experts know, since they tend to overestimate the amount of shared background knowledge.<sup>53</sup> They are thus a common source of unintentional stupefying. They are not the only source, though. Take someone who has bought into the quack's explanation, excitedly repeating it to their doctor friend hoping for it to be understood and confirmed. Here the friend may be stupefied (though hopefully they would not be), but not intentionally so.

### 3.2 Why Addressees are Stupefied

We've seen that there are a number of reasons speakers make potentially stupefying assertions, intentionally or otherwise. Why, though, do addressees allow themselves to be stupefied? Why go along with it?

In a personal conversation, when someone makes an assertion, here are the options the hearer has:



Besides confirming, denying, or raising some question about the utterance without confirming or denying it, one can of course do all sorts of things. Blankly stare into space, assert something totally unrelated, scream and run away, or stop, drop, and roll. These are all classed under the option *OTHER*, and are usually not options on an addressee's radar in a normal conversation, and so are typically ignored.<sup>54</sup> What we want to know is why, in a case in which something has been asserted which they don't

<sup>53</sup>A cognitive bias known as the curse of knowledge. See Camerer, Loewenstein, and Weber (1989), Birch and Bloom (2007), and Birch, Brosseau-Liard, et al. (2017).

<sup>54</sup>Why these and other options are usually excluded from consideration is an interesting question, but not one I will take up here.

understand, an addressee would pick CONFIRM, as opposed to DENY or QUESTION.

It's pretty clear why one usually wouldn't deny the utterance: since one doesn't understand it, it's unlikely one will be able to justify acceptance of the negation of what was asserted. And given the sort of challenge that denial typically involves to the original speaker, it's likely that this justification will be demanded. The same explanation goes, more or less, for raising some question that's supposed to bear on the assertion's truth. One's lack of understanding makes it hard to know what questions would make sense to raise, and the answers to them will likely just put one in a new stupor, so asking them doesn't accomplish much anyways.

What's more interesting is why, when faced with some assertion one doesn't understand, one doesn't simply ask for clarification about what the utterance means. Why not respond to the incomprehensible assertions with something like (13)?

- (13) a. I don't understand, could you explain that in other terms?  
b. What's an ionizing process/pizer/ferrite core/etc.?

Such requests for clarification are often the sensible thing to do, and indeed they are often made.<sup>55</sup> But sometimes addressees don't ask them, instead allowing themselves to be stupefied. They need not be irrational on this account, however, because clarification requests can have costs. The most common and important clarification costs, I think, stem from three facts: (i) they take time, (ii) they express ignorance, and (iii) they sometimes appear to challenge authority.

Regarding (i), it can take quite a while to explain to someone all they need to know in order to understand the relevant terms. In extreme cases—suppose someone has just used some arcane technical vocabulary—it can take years to acquire the concepts needed to understand. But even when it would take just a few minutes or less to figure out what's being talked about, this can still be a cost: one could be doing more useful things and so, one might worry, could one's interlocutor. If it doesn't seem important that one understand, allowing oneself to be stupefied may seem to be the best option.

As for (ii), we've already seen this cited as an explanation of stupefying:

---

<sup>55</sup>See Purver, Ginzburg, and Healey (2003) and Ginzburg (2012, Ch. 6).

Mrs. B., recall, worried that Dr. A. would have thought her to be stupid had she asked what was necessary to avoid being stupefied. This is a common reason that addressees allow themselves to be stupefied: by pretending to understand, one saves intellectual face.<sup>56</sup>

Moreover, in plenty of cases, ignorance can imply other failures which we wish to keep hidden. It is often one's responsibility to understand certain things because one is a citizen, say, or a parent, or a doctor, or because one has promised to read and think about something. Displaying ignorance by asking for clarification, then, can also involve the cost of revealing that one has failed to fulfill certain obligations.

Finally, our identities are often tied up with what we know, particularly with our knowledge of language.<sup>57</sup> One may worry that asking someone to tell them what *pizer* means may out them as an outsider (or a *dingbatter*, as it would be put in Ocracoke Brogue), or that asking what *supervenience* means would show one not to be a bona fide analytic philosopher. Allowing oneself to be stupefied is sometimes a way one tries to fit in.

Another potential social cost of clarification questions is (iii): that they can challenge authority, or at least appear to do so. We're all familiar with the aggressive "But what does that even mean?" clarification request meant less to find out what the speaker meant than to show everyone that they were talking nonsense. Clarification questions, even ones not meant in this way, can come across as challenges to the speaker, which can be costly.

One way it can be costly is when the speaker is, in one way or another, more powerful than the one asking the clarification question. If their authority appears to be challenged, they or their associates might impose some social or even physical cost on the questioner.

Another way the appearance of challenge can be costly is through the harmful effects it has on the person who made the original assertion. We can't always explain what we mean, even if we have said something perfectly meaningful. Nevertheless, someone's not being able to explain what they meant can make them appear incompetent, reducing their

---

<sup>56</sup>Sometimes, allowing oneself to be stupefied ends up displaying one's ignorance, as with Sokal's famous hoax (Sokal and Bricmont 1998, Ch. 1). An exception, I think, which proves the rule.

<sup>57</sup>The relationship between language and identity is core topic in sociolinguistics. See Mendoza-Denton (2002) and Tabouret-Keller (1997).

credibility. Clarification questions' ability to undermine credibility means they can exacerbate what Miranda Fricker (2007, p. 1) calls testimonial injustice, which occurs "when prejudice causes a hearer to give a deflated level of credibility to a speaker's word".<sup>58</sup> Triggering or worsening testimonial injustice is sometimes a cost of clarification questions. Stupefied acceptance is a way to avoid this cost.

How costly (i), (ii), and (iii) end up being varies a lot depending on the context—how pressed one is for time, what one is expected to know, the social positions of speaker and addressee, and so on. But it's not hard to see how in some cases they make requesting clarification into an unattractive option when compared with allowing oneself to be stupefied.

This is especially so when the costs of confirming are low, as they often are. One obvious way for the cost of confirming to be low is if not much rides on what ends up becoming joint discourse commitments, like when the conversational tone is not one of belief. If people are just chatting to pass the time and achieve various social goals unrelated to a conversation's content, there's not much harm in explicitly accepting some assertion which one doesn't understand. Hence the ubiquity of smiling and nodding at cocktail parties.

Another way for (expected) confirmation cost to be low is when what matters is not anything that depends on the hearer understanding, but just depends on whether what has been said is true. Sometimes the hearer has good evidence—and indeed can know—that what the speaker uttered is true, even without understanding.<sup>59</sup> So, for example, if my doctor tells me

---

<sup>58</sup>Indeed, it can do so in a particularly pernicious way, harnessing the other kind of epistemic injustice Fricker discusses—hermeneutical injustice—which occurs when "a gap in collective interpretive resources puts someone at an unfair disadvantage when it comes to making sense of their social experiences". This gap makes it more likely that the hermeneutically marginalized will describe their experiences in ways that others do not and cannot understand (Fricker 2007, p. 159). If responding to this with clarification questions lowers their credibility in a way at least partly do to prejudice, then they can be seen as a way of transforming hermeneutical into testimonial injustice.

<sup>59</sup>How to tell when one can trust an expert is an important question in social epistemology. See, for example, Goldman (2001). Interestingly, citing a suggestion by Carol Caraway, Goldman distinguishes between semantically and epistemically esoteric statements, where "[s]emantically esoteric statements are ones that a novice cannot assess because he does not even *understand* them; typically, they utilize a technical vocabulary he has not mastered. Epistemically esoteric statements are statements that the novice understands but still cannot assess for truth-value." (Goldman 2001, p. 94, n. 10). Semantically esoteric statements, of course, are just the sort of statements involved

that levothyroxine will “diffuse into the cell nucleus and bind to thyroid receptor proteins attached to DNA”, activating “gene transcription and synthesis of messenger RNA and cytoplasmic proteins” and that this is why she thinks it will help with some symptoms I have, it really makes no difference whether I understood what she said, so long as she’s right about it and about its connection to my symptoms.<sup>60</sup> And given her profession, I trust that she is right about this, so it seems like a low cost to just accept what she has said without understanding.<sup>61</sup> For similar reasons, it’s completely understandable how Eben Byers was stupefied into drinking Radithor.

### 3.3 What is to be Done?

So far in this section we’ve seen some reasons why speakers might say what their addressees won’t understand and why addressees sometimes accept what they say anyways. Sometimes speakers stupefy out of greed, malice, or carelessness, but not always. Sometimes addressees are stupefied out of vanity or obsequiousness, but not always. Stupefying can be a means of short-circuiting and manipulation for speakers and can be irrational for hearers, but it can also result from innocent, cooperative, and reasonable conversational moves. Not all stupefying is bad. We should probably not try to eliminate stupefying, then, but rather be wary of it. I’ll conclude by speculatively suggesting that some important goods are unattainable without stupefying. This would make any attempt to avoid stupefying altogether a serious mistake.

When we learn from others speaking with us—a central means of learning for humans—we usually do so by at least provisionally accepting

---

in stupefying.

<sup>60</sup>[http://www.accessdata.fda.gov/drugsatfda\\_docs/label/2009/021292s0021bl.pdf](http://www.accessdata.fda.gov/drugsatfda_docs/label/2009/021292s0021bl.pdf)

<sup>61</sup>One option which is often the best ones in situations like this is to confirm in the hedged way discussed in §2.3, accepting, but explicitly flagging one’s failure to understand. This has one of the costs of requesting clarification—it does express ignorance—but it can, if done in the right way, avoid the others. And by not giving a false impression of understanding, it can have some advantages over a more straightforward confirmation. For example, if the speaker thinks it is important that I understand, they can take the time to explain it. It’s also more clear to observers that any responsibility for defending it falls to the speaker.



what they say.<sup>62</sup> Some learning involves eliminating possibilities which we already realize, if only implicitly, are possibilities. This is very well modeled by the familiar Stalnakerian picture of conversation. Other learning, though, involves conceptual expansion, adding to our stock of what we consider possibilities.<sup>63</sup> We emerge from it not, primarily, with a collection of new facts about how the world is, but rather new ways of thinking about how the world might be. These new ways of thinking are not, generally, easily definable from what one already understands. So if a teacher aims to impart this kind of knowledge through speech, they will typically do so by saying what we don't yet understand. And to learn the new way of thinking from them, we provisionally accept what they say without understanding it, hoping that as we follow them and try thinking these thoughts on our own, things will begin to fall into place, and we gradually come to understand what we have already accepted. As with other kinds of learning, we learn new ways to think by doing. Learning from others in this way will involve stupefying.<sup>64</sup>

In one of Malebranche's dialogues, Theodore's pupil Aristes finally accepts something Theodore has been saying and restates the view back to him. "Do you conceive all this quite distinctly?" Theodore asks, "Are you quite convinced of it by the reasons I have given you and by your own reflections?" Aristes admits he was stupefied:

ARISTES. . . . It seems that you excite sensations in me instead

---

<sup>62</sup>On the centrality of testimony in our ancestors' way of life and its impact on the development of their minds and language, see Csibra and Gergely (2009), Sterelny (2012), Tomasello (2008), Tomasello (2014), and Mercier and Sperber (2017). On the crucial role of testimony in children's learning, see Harris (2012). There is also a large literature, of course, on the epistemology of testimony; see Lackey and E. Sosa (2006) for an entry point.

<sup>63</sup>For an interesting study of the kind of learning I have in mind, particularly in children, see Carey (2009). Her account of how it works, which she calls Quinian Bootstrapping, involves reliance on a kind of placeholder thought, to which content is added as it is used. A similar phenomenon is of interest to philosophers of science—see Strevens (2012) on what he calls 'introjection'. The epistemology of these sorts of transitions is, in my view, much understudied. But for some helpful discussions of epistemically evaluating concepts and conceptual schemes, see Carr (2015) and Pérez Carballo (forthcoming).

<sup>64</sup>This in turn means that, at least for a certain kind of learning by testimony, understanding of what the other said is not a precondition for success, but rather the result of it. This contrasts with the usual assumption that understanding is a precondition for testimony's epistemic success. For a very different sort of challenge to the view that testimonial success requires understanding, see Peet (2018).

of producing clear ideas. I am using your language. In all honesty, I do not understand everything you are telling me. I see it, and a moment later I no longer see it. For I still only glimpse it. It seems to me you are right, but I am not understanding you too well.

Malebranche (1688/1997, §III.viii)

As we're learning new ways to think, we often only barely glimpse what someone is saying. To come to understand it, we accept it anyways, using their language until it becomes our own. If this is right, we need to live with stupefying. There are times when we should resist the temptation to stupefy and be stupefied, but other times it is necessary for imparting and acquiring new ways of thinking.

## References

- Ågotnes, Thomas and Hans van Ditmarsch (2011). "What will they say?—Public Announcement Games". In: *Synthese* 179, Supplement 1, pp. 57–85.
- Akerlof, George A. and Robert J. Shiller (2015). *Phishing for Phools: The Economics of Manipulation and Deception*. Princeton University Press.
- Almotahari, Mahrad and Ephraim Glick (2010). "Context, Content, and Epistemic Transparency". In: *Mind* 119.476, pp. 1067–1086.
- Anderbois, Scott, Adrian Brasoveanu, and Robert Henderson (2015). "At-issue proposals and appositive impositions in discourse". In: *Journal of Semantics* 32, pp. 93–138.
- Asher, Nicholas and Alex Lascarides (2013). "Strategic conversation". In: *Semantics and Pragmatics* 6.
- Asher, Nicholas and Soumya Paul (2018). "Strategic Conversations Under Imperfect Information: Epistemic Message Exchange Games". In: *Journal of Logic, Language, and Information* 27, pp. 343–385.
- Beaver, David and Jason Stanley (forthcoming). *Hustle: The Politics of Language*. Princeton University Press.
- (2019). "Toward a Non-Ideal Philosophy of Language". In: 39.2, pp. 501–545.
- Birch, Susan A.J. and Paul Bloom (2007). "The curse of knowledge in reasoning about false beliefs". In: *Psychological Science* 18.5, pp. 382–386.
- Birch, Susan A.J., Patricia E Brosseau-Liard, et al. (2017). "A 'curse of knowledge' in the absence of knowledge? People misattribute fluency when judging how common knowledge is among their peers". In: *Cognition* 166, pp. 447–458.
- Bledin, Justin and Kyle Rawlins (forthcoming). "Resistance and Resolution: Attentional Dynamics in Discourse". In: *Journal of Semantics*.
- (2016). "Epistemic resistance moves". In: *Proceedings of SALT*, pp. 620–640.
- Buchanan, Ray (2010). "A Puzzle about Meaning and Communication". In: *Noûs* 44.2, pp. 340–371.
- Burge, Tyler (1979). "Individualism and the mental". In: *Midwest Studies in Philosophy* 4.1, pp. 73–122.

- Camerer, Colin, George Loewenstein, and Martin Weber (1989). "The Curse of Knowledge in Economic Settings: An Experimental Analysis". In: *Journal of Political Economy* 97.5, pp. 1232–1254.
- Camp, Elisabeth (2018). "Insinuation, Common Ground, and the Conversational Record". In: *New Work on Speech Acts*. Ed. by Daniel Fogal, Daniel W. Harris, and Matt Moss. Oxford University Press, pp. 40–66.
- Cappelen, Herman and Josh Dever (2019). *Bad Language*. Oxford University Press.
- Cappelen, Herman and John Hawthorne (2009). *Relativism and Monadic Truth*. Oxford University Press.
- Carey, Susan (2009). *The Origin of Concepts*. Oxford University Press.
- Carr, Jennifer (2015). "Epistemic Expansions". In: *Res Philosophica* 92.2, pp. 217–236.
- Ciardelli, Ivano, Jereon Groenendijk, and Floris Roelofsen (2009). "Attention! *Might* in Inquisitive Semantics". In: *Proceedings of Semantics and Linguistic Theory* 19, pp. 91–108.
- Clark, Claudia (1997). *Radium Girls: Women and Industrial Health Reform, 1910–1935*. University of North Carolina Press.
- Clark, Herbert H. and Catherine R. Marshall (1981). "Definite Reference and Mutual Knowledge". In: *Elements of Discourse Understanding*. Ed. by Aravind Joshi, Bonnie Webber, and Ivan Sag. Cambridge University Press, pp. 10–63.
- Crone, Phil (2017). "Redundancy and Awareness in Discourse". PhD thesis. Stanford University.
- (2018). "Assertions of Clarity & Raising Awareness". In: *Journal of Semantics* 36.1, pp. 53–97.
- Csibra, Gergely and György Gergely (2009). "Natural pedagogy". In: *Trends in Cognitive Science* 13.4, pp. 148–153.
- Dummett, Michael (1973). *Frege: Philosophy of Language*. London: Duckworth.
- Fagin, Ronald and Joseph Y. Halpern (1988). "Belief, Awareness, and Limited Reasoning". In: *Artificial Intelligence* 34, pp. 39–76.
- Farkas, Donka F. and Kim B. Bruce (2010). "On Reacting to Assertions and Polar Questions". In: *Journal of Semantics* 27 (1), pp. 81–118.

- Farkas, Donka F. and Floris Roelofsen (2017). "Division of Labor in the Interpretation of Declaratives and Interrogative". In: *Journal of Semantics* 34, pp. 237–289.
- Ferrante, Elena (2014). *Those Who Leave and Those Who Stay*. Translated by Ann Goldstein. Europa Editions.
- Fine, Kit (2017). "A Theory of Truth-Conditional Content I: Conjunction, Disjunction and Negation". In: *Journal of Philosophical Logic* 46, pp. 625–674.
- Fletcher, John (1976). "Realities of Patient Consent to Medical Research". In: *Biomedical Ethics and the Law*. Ed. by James M. Humber and Robert F. Almeder. Plenum Press, pp. 261–276.
- Franke, Michael (2014). "Pragmatic Reasoning About Unawareness". In: *Erkenntnis* 79, pp. 729–767.
- Franke, Michael and Tikitou de Jager (2011). "Now That You Mention It: Awareness Dynamics in Discourse and Decisions". In: *Language, Games, and Evolution: Trends in Current Research on Language and Game Theory*. Ed. by Anton Benz et al. Springer, pp. 60–91.
- Frankfurt, Harry G. (1982). "On Bullshit". In: *Raritan* 6.2. Reprinted in Frankfurt (1998), pp. 117–133).
- (1998). *The Importance of What We Care About*. Cambridge University Press.
- Fricker, Miranda (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.
- Fritz, Peter and Harvey Lederman (2015). "Standard State Models of Unawareness". In: *Proceedings of TARK XV: Fifteenth Conference on Theoretical Aspects of Rationality and Knowledge*, pp. 163–172.
- Ginzburg, Jonathan (1996). "Dynamics and the Semantics of Dialogue". In: *Logic, language, and computation*. Ed. by Jerry Seligman and Dag Westersåhl. Vol. 1. CSLI Lecture Notes 58. Center for the Study of Language & Information, Stanford University.
- (2012). *The Interactive Stance*. Oxford University Press.
- Goldman, Alvin I. (2001). "Experts: Which Ones Should You Trust?" In: *Philosophy and Phenomenological Research* 63.1, pp. 85–110.
- Grice, H. Paul (1967). "Logic and Conversation". Printed in Grice (1989, pp. 41–58).

- Grice, H. Paul (1989). *Studies in the Way of Words*. Cambridge, Massachusetts: Harvard University Press.
- Halpern, Joseph Y. and Evan Piermont (2019). "Partial Awareness". In: *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*.
- Halpern, Joseph Y. and Leandro C. Rêgo (2014). "Extensive games with possibly unaware players". In: *Mathematical Social Sciences* 70, pp. 42–58.
- Hamblin, Charles (1971). "Mathematical models of dialogue". In: *Theoria* 37, pp. 130–155.
- Harris, Paul L. (2012). *Trusting What You're Told: How Children Learn from Others*. The Belknap Press of Harvard University Press.
- Haslanger, Sally (2011). "Ideology, Generics, and Common Ground". In: *Feminist Metaphysics: Explorations in the Ontology of Sex, Gender and the Self*. Ed. by Charlotte Witt. Springer, pp. 179–207. Reprinted in Haslanger (2012), pp. 446–477).
- (2012). *Resisting Reality: Social Construction and Social Critique*. Oxford University Press.
- Hawthorne, John and Ofra Magidor (2009). "Assertion, Context, and Epistemic Accessibility". In: *Mind* 118.470, pp. 377–397.
- Heifetz, Aviad, Martin Meier, and Burkhard C. Schipper (2013). "Dynamic unawareness and rationalizable behavior". In: *Games and Economic Behavior* 81, pp. 50–68.
- Henderson, Robert and Elin McCready (forthcoming). "Dogwhistles and the at-issue/not-at-issue distinction". In: *Secondary Information: The linguistics of side issues*. Ed. by Daniel Gutzmann and Katharina Turgay. Brill Academic Publishers.
- Kaplan, David (1989). "Demonstrativeness: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals". In: *Themes from Kaplan*. Ed. by Joseph Almog, John Perry, and Howard Wettstein. Oxford University Press, pp. 481–564.
- Karttunen, Lauri (1974). "Presuppositions and Linguistic Context". In: *Theoretical Linguistics* 1, pp. 181–194.
- Khoo, Justin and Rachel Sterken, eds. (forthcoming). *The Routledge Handbook of Social and Political Philosophy of Language*. Routledge.
- King, Jeffrey C. (2007). *The Nature and Structure of Content*. Oxford University Press.

- King, Jeffrey C. (2018). "Strong Contextual Felicity and Felicitous Underspecification". In: *Philosophy and Phenomenological Research* 97.3, pp. 631–657.
- Lackey, Jennifer and Ernest Sosa, eds. (2006). *The Epistemology of Testimony*. Oxford University Press.
- Langton, Rae (forthcoming). *Accommodating Injustice: The John Locke Lectures 2015*. Oxford University Press.
- (2018a). "Blocking as Counter-Speech". In: *New Work on Speech Acts*. Ed. by Daniel Fogal, Daniel W. Harris, and Matt Moss. Oxford University Press, pp. 144–164.
- (2018b). "The Authority of Hate Speech". In: *Oxford Studies in Philosophy of Law*. Ed. by John Gardner, Leslie Green, and Brian Leiter. Vol. 3, pp. 123–152.
- Lederman, Harvey (2018). "Uncommon Knowledge". In: *Mind* 127 (508), pp. 1069–1105.
- Lepore, Ernie and Matthew Stone (2015). *Imagination and Convention: Distinguishing Grammar and Inference in Language*. Oxford University Press.
- Lewis, David (1969). *Convention*. Harvard University Press.
- (1979). "Scorekeeping in a language game". In: *Journal of Philosophical Logic* 8, pp. 339–359. Reprinted in Lewis (1983, pp. 233–249).
- (1980). "Index, Context and Content". In: *Philosophy and Grammar*. Ed. by S. Kanger and S. Oshman. Dordrecht: Reidel, pp. 79–100. Reprinted in Lewis (1998, pp. 21–44).
- (1983). *Philosophical Papers*. Vol. I. Oxford: Oxford University Press.
- (1998). *Papers in Philosophical Logic*. Cambridge University Press.
- MacFarlane, John (2011). "What Is Assertion?" In: *Assertion: New Philosophical Essays*. Ed. by Jessica Brown and Herman Cappelen. Oxford University Press, pp. 79–96.
- Malebranche, Nicolas (1688/1997). *Dialogues on Metaphysics and Religion*. Ed. by Nicholas Jolley. Trans. by David Scott. Cambridge University Press.
- McGowan, Mary Kate (2004). "Conversational Exercises: Something Else We Do with Our Words". In: *Linguistics and Philosophy* 27.1, pp. 93–111.

- McGowan, Mary Kate (2018). "On Covert Exercitives: Speech and the Social World". In: *New Work on Speech Acts*. Ed. by Daniel Fogal, Daniel W. Harris, and Matt Moss, pp. 185–201.
- (2019). *Just Words: On Speech and Hidden Harm*. Oxford University Press.
- Mendoza-Denton, Norma (2002). "Language and Identity". In: *The Handbook of Language Variation and Change*. Ed. by J. K. Chambers, Peter Trudgill, and Natalie Schilling-Estes, pp. 475–499.
- Mercier, Hugo and Dan Sperber (2017). *The Enigma of Reason*. Harvard University Press.
- Mole, Christopher (2011). *Attention Is Cognitive Unison: An Essay in Philosophical Psychology*. Oxford University Press.
- (2017). "Attention". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Fall 2017. Metaphysics Research Lab, Stanford University.
- Morris, Charles Evans (1926). *Modern Rejuvenation Methods*. 249 West 34th St., New York: Scientific Medical Publishing Co.
- Moss, Sarah (2018). *Probabilistic Knowledge*. Oxford University Press.
- Murray, Sarah E. and William B. Starr (forthcoming). "The Structure of Communicative Acts". In: *Linguistics and Philosophy*.
- Ninan, Dilip (2010). "Semantics and the objects of assertion". In: *Linguistics and Philosophy* 33, pp. 355–380.
- Peet, Andrew (2018). "Testimonial Knowledge Without Knowledge of What is Said". In: *Pacific Philosophical Quarterly* 99, pp. 65–81.
- Pennycook, Gordon et al. (2015). "On the reception and detection of pseudo-profound bullshit". In: *Judgement and Decision Making* 10.6, pp. 549–563.
- Pérez Carballo, Alejandro (forthcoming). "Conceptual Evaluation: Epistemic". In: *Conceptual Engineering and Conceptual Ethics*. Ed. by Herman Cappelen Alexis Burgess and David Plunkett. Oxford University Press.
- Potts, Christopher (2005). *The Logic of Conventional Implicatures*. Oxford Studies in Theoretical Linguistics. Oxford: Oxford University Press.
- Purver, Matthew, Jonathan Ginzburg, and Patrick Healey (2003). "On the means for clarification in dialogue". In: *Current and new directions in discourse and dialogue*. Springer, pp. 235–255.
- Rabern, Brian (2012). "Against the Identification of Assertoric Content with Compositional Value". In: *Synthese* 189.1, pp. 75–96.



- “Radium Poison in Tonic Kills Eben M. Byers” (1932). In: *Chicago Daily Tribune* (Friday, April 1), pp. 1, 16.
- Recanati, François (1997). “Can We Believe What We Do Not Understand”. In: *Mind & Language* 12.1, pp. 84–100.
- Roberts, Craige (1996/2012). “Information structure in discourse: Towards an integrated formal theory of pragmatics”. In: *Semantics and Pragmatics* 5.
- Roberts, Craige et al. (2009). “Presuppositions, conventional implicature, and beyond: a unified account of projection”. In: *Proceedings New Directions in the Theory of Presupposition (ESSLLI)*. Ed. by Nathan Klinedinst and Daniel Rothschild, pp. 1–15.
- Roelofsen, Floris (2013). “A bare bones attentive semantics for *might*”. In: *The dynamic, inquisitive, and visionary life of  $\phi$ ,  $?\phi$ , and  $\diamond\phi$ : a festschrift for Jereon Groenendijk, Martin Stokhof, and Frank Veltman*. Ed. by Maria Aloni, Michael Franke, and Floris Roelofsen. Amsterdam: Institute for Logic, Language, and Computation, pp. 190–215.
- Saul, Jennifer (2018). “Dogwhistles, Political Manipulation, and Philosophy of Language”. In: *New Work on Speech Acts*. Ed. by Daniel Fogal, Daniel W. Harris, and Matt Moss, pp. 360–384.
- Schiffer, Stephen (1972). *Meaning*. Oxford University Press.
- Schipper, Burkhard C. (2015). “Awareness”. In: *Handbook of Epistemic Logic*. Ed. by Hans van Ditmarsch et al. College Publications, pp. 77–146.
- Smith, Adam (1790). *The Theory of Moral Sentiments*. 6th ed. London: A. Strahan and T. Cadell.
- Sokal, Alan and Jean Bricmont (1998). *Fashionable Nonsense: Postmodern Intellectuals’ Abuse of Science*. New York: Picador.
- Sosa, David, ed. (2018). *Bad Words: Philosophical Perspectives on Slurs*. Oxford University Press.
- Sperber, Dan (2010). “The Guru Effect”. In: *Review of Philosophical Psychology* 1, pp. 583–592.
- Stalnaker, Robert (1970). “Pragmatics”. In: *Synthese* 22. Reprinted in Stalnaker (1999), pp. 31–46).
- (1974). “Pragmatic Presuppositions”. In: *Semantics and Philosophy*. Ed. by M. Munitz and P. Unger. New York University Press, pp. 197–214. Reprinted in Stalnaker (1999).

- Stalnaker, Robert (1978). "Assertion". In: *Syntax and Semantics* 9, pp. 315–332. Reprinted in Stalnaker (1999).
- (1984). *Inquiry*. MIT Press.
- (1998). "On the Representation of Context". In: *Journal of Logic, Language, and Information* 7. Reprinted in Stalnaker (1999, pp. 96–113).
- (1999). *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford University Press.
- (2002). "Common Ground". In: *Linguistics and Philosophy* 25.5/6, pp. 701–721.
- (2004). "Assertion Revisited: On the Interpretation of Two-Dimensional Modal Semantics". In: *Philosophical Studies* 118, pp. 299–322.
- (2009). "On Hawthorne and Magidor on Assertion, Context, and Epistemic Accessibility". In: *Mind* 118.470, pp. 399–409.
- (2014). *Context*. Oxford University Press.
- (2018). "David Lewis on Context". In: *The Science of Meaning: Essays on the Metatheory of Natural Language Semantics*. Ed. by Derek Ball and Brian Rabern. Oxford University Press, pp. 98–112.
- Stanley, Jason (1997). "Rigidity and Content". In: ed. by Richard Kimberly Heck, pp. 131–156.
- (2015). *How Propaganda Works*. Princeton University Press.
- Sterelny, Kim (2012). *The Evolved Apprentice: How Evolution Made Humans Unique*. The MIT Press.
- Stojnić, Una (2017). "On the Connection between Semantic Content and the Objects of Assertion". In: *Philosophical Topics* 45.2, pp. 163–180.
- Stokke, Andreas (2018). *Lying and Insincerity*. Oxford University Press.
- Strevens, Michael (2012). "Theoretical terms without analytic truths". In: *Philosophical Studies* 160.1, pp. 167–190.
- Syrett, Kristen and Todor Koev (2014). "Experimental evidence for the truth conditional contribution and shifting information status of appositives". In: *Journal of Semantics* 32.3, pp. 525–577.
- Tabouret-Keller, Andrée (1997). "Language and Identity". In: *The Handbook of Sociolinguistics*. Ed. by Florian Coulmas, pp. 315–326.
- Tirrell, Lynne (2018). "Toxic Speech: Inoculation and Antidotes". In: *The Southern Journal of Philosophy* 56, pp. 116–144.
- Tomasello, Michael (2008). *Origins of Human Communication*. The MIT Press.

- Tomasello, Michael (2014). *A Natural History of Human Thinking*. Harvard University Press.
- Van Benthem, Johan (2011). *Logical Dynamics of Information and Interaction*. Cambridge University Press.
- Von Fintel, Kai (2008). "What is Presupposition Accommodation, Again?" In: *Philosophical Perspectives* 22, pp. 137–170.
- Watzl, Sebastian (2017). *Structuring Mind: The Nature of Attention and How it Shapes Consciousness*. Oxford University Press.
- Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford University Press.
- Winslow, Ron (1990). "The Radium water Worked Fine Until His Jaw Came Off". In: *The Wall Street Journal* (August 1), A1, A5.
- Wolfram, Walt and Jeffrey Reaser (2014). *Talkin' Tar Heel: How Our Voices Tell the Story of North Carolina*. University of North Carolina Press.
- Yablo, Stephen (2014). *Aboutness*. Princeton University Press.
- Yalcin, Seth (2007). "Epistemic Modals". In: *Mind* 116.464, pp. 983–1026.
- (2012). "Context Probabilism". In: *Logic, Language, and Meaning: 18th Amsterdam Colloquium Revised Selected Papers*. Ed. by Maria Aloni et al., pp. 12–21.
- (2016). "Belief as Question-Sensitive". In: *Philosophy and Phenomenological Research*, pp. 1–25.