

Truth and imprecision

Josh Armstrong 

Department of Philosophy, University of California at Los Angeles, Los Angeles, California, USA

Correspondence

Josh Armstrong, Department of Philosophy, University of California at Los Angeles, 390 Portola Plaza, Los Angeles, CA 90095, USA.
Email: jarmstrong@humnet.ucla.edu

Abstract

Our ordinary assertions are often imprecise, insofar as the way we represent things as being only approximates how things are in the actual world. The phenomenon of *assertoric imprecision* raises a challenge to standard accounts of both the norm of assertion and the connection between semantics and the objects of assertion. After clarifying these problems in detail, I develop a framework for resolving them. Specifically, I argue that the phenomenon of assertoric imprecision motivates a rejection of the widely held belief that a semantic theory for a language associates a single semantic value with each of the simple and complex expressions of that language, relative to the contexts in which they occur. Instead, I propose that we adopt a framework I call *semantic pluralism*.

1 | INTRODUCTION

Alf and Bea are standing in the library lobby. Alf asks Bea for the time. Bea looks at her watch, and replies “It’s three o’clock.” Alf thanks her and goes on his way. Alf and Bea have had a successful conversational exchange: Alf requested some information that he lacked, Bea acquired the information in question and uttered a sentence that served to communicate it. Alf understood what information the sentence was intended to express, trusted Bea, and updated with that information accordingly.

One of the aims of the philosophy of language is to provide a systematic account of the practice whereby languages are used for the purposes of interpersonal communication. When we reflect on humdrum cases of interpersonal communication, like that of Alf and Bea above, various kinds of regularities seem to emerge. By and large, speakers perform speech acts with the intention to provide accurate information and, by and large, hearers expect this from speakers. In addition, speakers by and large perform speech acts with the intention to be understood and hearers likewise expect that speakers will speak in a way that can be understood as intended. The fact that these two kinds of regularities obtain—and, indeed, may well be matters of common knowledge among those involved—seems to provide us with the makings of a theory that explains that remarkable success with which languages are used for the purposes of interpersonal communication.

There is, however, a problem. Speakers often speak loosely: the information they provide often only approximates how things are in the actual world. In our humdrum case above, let us suppose that it was 3:03. Indeed, let us suppose that Bea's watch read 3:03. It thus seems that it was not actually 3:00, as Bea's assertion represented things as being. For reasons that will emerge I'll call this a case of *assertoric imprecision*.

There can be no denying that speakers do often speak loosely in this way. Our ordinary practices of linguistic communication would be strained if they did not: speakers would be far more pedantic in what they said, hearers far less interested in what was said, and the efficiency of linguistic communication would be threatened. However, what has been matter of theoretical controversy is the implications that cases of assertoric imprecision have for the nature of semantic theory—in particular, for the relationship between semantics, truth, and communication.

In recent years, a number of theorists have taken cases of assertoric impression to suggest serious revisions to the platitudes with which we began. Some theorists have used the problem to motivate a rejection of standard alethic norms on assertion. Wilson and Sperber (2002) have taken such cases to suggest that assertion is not governed by a norm of truth or knowledge but instead a norm of *relevance* (c.f. Sperber and Wilson [1986]). Other theorists, for example, Lasersohn (1999), have taken the data to suggest that the semantic content of a sentence is rarely the primary information a speaker intends to convey in asserting that sentence. Instead, alethic norms on assertion govern propositions that are pragmatically conveyed over and above the literal meanings of speakers' words.

My aim in this paper is develop and defend an account of assertoric imprecision according to which there is an alethic norm on assertion and that norm governs the literal meanings of speakers' words. I will argue that what stands to be revised is a widespread but tacit assumption concerning the relationship between conversational participants and the semantic systems that they use for the purposes of interpersonal communication. I'll label this assumption *semantic uniqueness* because it holds that linguistic expressions are conventionally associated with unique semantic contents in the contexts in which they are used. Semantic uniqueness should be rejected, and, in its place, we should adopt a version of what I'll call *semantic pluralism*. Semantic pluralism holds that linguistic expressions can come to be associated with multiple semantic contents in the contexts in which they are used, and, as a consequence, a speaker may literally express multiple propositions in asserting a single sentence.

This pluralist approach to semantic content allows us to see how speakers could be *literally truthful* with their assertions while still being quite imprecise, and this, I will argue, is crucial for resolving the problem. More generally, semantic pluralism allows us to take the phenomena of assertoric imprecision seriously without rejecting either a constitutive connection between assertion and truth or a tight connection between literal meaning and assertoric content.

The plan for the paper is as follows: in Section 2, I show how the problem of assertoric imprecision arises within one particular version of a standard theory of assertion, adapted from the work of Robert Stalnaker and David Lewis, and I identify some desiderata that any adequate account of assertoric imprecision should satisfy. In Section 3, I consider the two aforementioned responses to the problem and argue that they should be rejected. In Section 4, I take some modest, initial steps toward articulating a pluralist approach to semantics. I close the paper, in Section 5, by working to show that semantic pluralism meets the desiderata for an account of assertoric imprecision.

2 | ASSERTION, TRUTH, AND SEMANTICS

In what follows, my focus shall be on *assertoric* uses of language. Assertions are intentional actions—actions performed with the goal of bringing about characteristic sorts of effects. I will assume a few

broad points about the nature of assertion that are more-or-less assumed by all parties to the dispute that I shall discuss. First, assertions are intentional actions that express *propositional contents*—entities that, whatever else they are, can be evaluated for truth or falsity relative to a state of the world or a point of evaluation more generally. Second, speakers intend to communicate the propositional contents of their assertions to the other members of a conversation and *communicative success is measured by the degree to which the members of a conversation coordinate on the propositional content(s) that a speaker intends to communicate*.

Of course, it would be nice to have a slightly more rigorous characterization of assertion. To this end, I will work with an abstract framework that models assertion in terms of the way it serves to update to the *common ground* of a conversation.¹ The common ground of a conversation is a body of information that is taken by the conversational participants to be mutually accepted for the purposes of a conversation; slightly more formally, the common ground is a set of propositions that is accepted by each conversational participant, and that is believed by each participant to likewise be accepted by each other member of the conversation.² The common ground of a conversation can, in this sense, be understood to be a kind of mutually accessible public record of a conversation.

Within this model, assertions are distinguished from other uses of language—such as questioning and commanding—in terms of the specific way they serve to eliminate *open possibilities* from consideration. These open possibilities are ways the world could be that are left open by what has been mutually accepted as true by the members of the conversation; roughly, open possibilities are possible worlds that the conversational participants are treating as live candidates for the actual world. The essential effect of an assertion is to update the common ground by eliminating open possibilities inconsistent with the content of that assertion. That is, in making an assertion the speaker is offering a proposal to the other members of the conversation that they narrow the set of open possibilities provided by the common ground to those in which the content of the assertion is true.³

Assertoric imprecision generates a problem given two widely accepted claims concerning assertion. The first concerns the connection between assertion and truth, and the second concerns the connection between assertoric content and semantics.

First, it is widely maintained that there is a constitutive connection between assertion and alethic notions like truth or knowledge. Although the literature on assertion is replete with differing versions of such a claim, I'll focus on the following version of the principle:

(TRU): As a default, one should only make an assertion provided that one takes that assertion to express a proposition which is true (in the world of evaluation).

The principle is not merely supposed to be a statistical generalization. Rather, (TRU) expresses a norm governing the practice of making assertions—something to which speakers are committed and to which audience members expect that speakers are committed. Indeed, the principle is said to be constitutive of assertion so that were speakers not committed to (TRU), they would not be making assertions.

In terms of the model of assertion involving the common ground, (TRU) amounts to the claim that there is a norm according to which speakers are to only assert contents which are presumed to narrow the live possibilities in context to those containing the actual world. This is to say, conversational

¹The model is taken from Stalnaker (1978) and Lewis (1979), who were in turn building on Grice (1957).

²As Stalnaker (1984) emphasizes, one can accept a proposition without believing it to be true—for instance, one might accept a proposition as true merely for the purposes of the conversation or inquiry.

³See Murray (2014) and Roberts (2018) for more discussion and for comparison with questions and commands.

participants expect that the actual world will be among the possibilities remaining when they update with the content of a speaker's assertion.

Second, and more controversially, many theorists have endorsed a tight connection between the content asserted by an utterance of a sentence in a context and the conventionally determined semantic content of that sentence in that context. Specifically, many theorists have adopted the following claim:

(SEM): As a default, the content expressed by an assertion of sentence *S* in context *C* is the conventionally determined semantic content of *S* in *C* (or the “literal meaning” of the sentence asserted).⁴

This claim is motivated by the familiar idea that the aim of a semantic theory of a language is to provide an account of the relation between sentences of a language and the information about the world that is conventionally communicated by utterances of those sentences (Kennedy & Stanley, 2009; Lewis, 1975; Stalnaker, 2014; among many others).

Since linguistic communication is both systematic and productive, it is generally assumed that adequate semantic theories should derive the truth-conditional content associated with sentences in context as a function of the denotation—or *semantic value*—of the basic expressions occurring in those sentences together with their manner of combination. The semantic content of a sentence, and the semantic value of an expression, at a context of utterance are then evaluated for truth and reference relative to a circumstance of evaluation (which I'll simply treat as a possible world). In terms of the common ground model of assertion, (SEM) amounts to the claim that the conversational participants have coordinated on a semantic procedure that—in conjunction with other cognitive abilities—allows them to determine which possibilities to eliminate on the basis of accepting a speaker's assertion of a sentence in a particular context of utterance.

Putting these two claims together, we can say that a semantic theory for a language elucidates what members of a group can use that language to say, and there is a norm according to which speakers only assert semantic contents which they take to be true. Of course, this way of understanding the aims of semantics has many detractors—including those who maintained that we must bifurcate the semantic values of sentences in context and the content that assertions of those sentences can be used to communicate.⁵ Addressing the concerns raised by skeptics of the traditional conception of the relationship between semantic value and assertoric content is an interesting and important project (cf. Stojnic, 2017a, 2017b), but it will not be my concern here. Instead, I'll treat the traditional conception of semantics as a working hypothesis in the discussion that follows.

2.1 | Assertoric imprecision

The foregoing account of the relation between semantics, truth and assertion looks to be at odds with the existence of cases of what I call *assertoric imprecision*. In cases of assertoric imprecision,

⁴Note that this principle is merely a default, so it can allow for assertoric content and semantic content to diverge in specific cases—for example, in cases of conversational implicatures, irony, satire, and the like.

⁵For some recent examples, see Carston (2002), Soames (2002, 2009), Ninan (2010), Rabern (2012), Yli-Vakkuri (2013), and Yalcin (2014) among many others.

speakers represent the world as being a certain way with their assertions that only approximates how things are in the actual world. Here are three characteristic examples:

- (1) “The bell rang at 6:00 p.m.” Known Fact: The bell rang at 6:03 p.m.
- (2) “I’m going to a conference in Chicago this weekend.” Known Fact: The conference is in Evanston, just outside of Chicago.
- (3) “8,000 people attended the rally today.” Known Fact: 7,998 people attended the rally.

In each case, it appears that the semantic content of the speakers' assertions—e.g., the literal meaning of their assertions—are not true given the facts which were known to obtain by the speakers in question. But this seems to be at odds either with (TRU) the claim that truth is the primary norm of assertion or (SEM) the claim that it is the semantic content of a speaker's assertion that we are evaluating for truth or falsity.

There is nothing particularly exceptional about (1)–(3): ordinary discourse is replete with assertions that equally well generate the problem. That speakers do assert sentences like (1)–(3) in the described circumstances should not be controversial. Neither should it be controversial that such utterances constitute bona fide assertions. What is at issue concerns whether the semantic content of the speakers' assertions—their literal meanings—are not true given the actual facts. Indeed, it is the primary aim of this paper to argue that, perhaps despite appearances, the semantic contents of sentences such as (1)–(3) are literally true in the contexts in which they were made. In this sense, I will argue that there is no problem of assertoric imprecision because the cases are perfectly consistent with (TRU) and (SEM). But it would be preferable if the debate did not simply turn on mere differences in intuitive truth-value judgments. For this reason, it will be worthwhile to consider some of the more theoretical reasons that seem to support the claim that assertoric imprecision involves cases of literal falsehoods. I will focus on two such considerations.

The first consideration centers on data involving what—borrowing a term introduced in a related context by DeRose (1995)—I'll call *abominable conjunctions*. Consider assertions of the following two sentences:

- (4) #“Although the bell rang at 6:00, it wasn't until 6:03 that it rang.”
- (5) #“I am going to a conference in Chicago this weekend; it is in Evanston, just outside of Chicago.”

There is something clearly defective with asserting (4) or (5). Indeed, Lasersohn (1999) has argued that these assertions express semantic contradictions: there is no context and world in which the two conjoined propositions could both be true.

But now contrast (4) and (5), with (4') and (5') that contain the presence of so-called *slack regulators* that serve to modify the content of the expressions with which they compositionally interact:

- (4') “At approximately 6:00 the bell rang. Indeed, it rang at 6:03.”
- (5') “I am going to a conference around Chicago this weekend; it is in Evanston, just outside of Chicago.”

These sentences seem perfectly assertable and clearly are not semantic contradictions. What these data have been taken to suggest is that if an assertion of (1) above could be literally true even if uttered at 3:03, we would expect that (4) should be both true and assertable, as we find with (4'). Since this is not what we find, there is reason to think what is semantically expressed by asserting sentences such (1)—(3) is literally false.

A second consideration centers on various kinds of *retraction data*. Consider the following dialogue:

Alf: 8,000 people attended the rally today. Bea: That's not true: there were only 7,988 people at the rally today.

Whether or not Bea is being unreasonable or pedantic in this exchange, a natural response is that she is correct: what Alf literally said was not true. Indeed, speakers often concede that they spoke falsely with their assertions in such cases. For example, a natural continuation of this dialogue is as follows:

Alf: 8,000 people attended the rally today. Bea: That's not true: there were only 7,988 people at the rally today. Alf: OK, fine, you're right: there were only 7,998 people at the rally today.

The question is why speakers would be so willing to retract their assertion, if what they said were not literally false? A straightforward answer to this question is that appearances are as they seem, and speakers are uttering literal falsehoods with their assertions in such cases.

An adequate treatment of assertoric imprecision should specify what modifications, if any, need to be made to (TRU) and (SEM); furthermore, it should do so in a way that addresses the data involving abominable conjunctions and retractions. But before going on to consider attempts to do this, it is perhaps worth first making clear what the problem of assertoric imprecision is not.

First, the problem of assertoric imprecision should be distinguished from the many problems raised by the vagueness displayed by expressions such as 'heap' or 'bald.' According to a standard view, the semantic values of vague expressions have *borderline cases*—entities that neither definitely fall within the extension of the term nor definitely do not fall within the extension of the terms—and, furthermore, there are *no sharp boundaries* between the respective extensions, anti-extensions, and borderline cases of the terms at issue. But cases of assertoric imprecision need not have borderline cases, and there may well be sharp boundaries between the extension and anti-extension of the semantic values of the terms involved. For example, the problem of imprecision generated by an assertion of (3) seems to remain if we are in an idealized situation in which there is no vagueness associated with terms "eight thousand," "people," or "attended." So the problem of assertoric imprecision would remain even if we assumed there was no semantic vagueness present, and each of the expressions occurring in the relevant sentences had perfectly determinate extensions with sharp boundaries.⁶ Semantic vagueness and semantic imprecision are closely related phenomena, but they should not be conflated with one another.

⁶If one adopts an epistemicist conception of vagueness, of the sort defended by Timothy Williamson, then the extensions of vague expressions will not have borderline cases, and will in fact determine sharp boundaries. But the epistemicist does maintain that we cannot—as a matter of principle—say what the sharp boundaries are for vague terms because these *boundaries are unknowable for us*. But in cases of assertoric imprecision, agents may be perfectly capable of knowing and even specifying the boundaries of the terms they are using. So even an epistemicist can accept the present distinction between vagueness and imprecision.

The problems raised by cases of assertoric imprecision should also be distinguished from cases of assertion involving *pretense*. It is a familiar point (Stalnaker, 1984) that speakers sometimes presuppose propositions for the purposes of a conversation or an inquiry that they believe or know to be false. For example, in a conversation with a family member about politics, a speaker may presuppose things that she does not believe to be true in order to focus the conversation on other things or simply to be polite. In cases like this, speakers operate under a pretense and assert propositions that they do not take to correspond to the actual facts—representing the world as they think other members of the conversation should take it to be rather than the way they themselves take the world to be. In contrast, speakers need not and often are not presupposing propositions they take to be false in cases of assertoric imprecision. Accordingly, the problem of assertoric imprecision cannot be sidestepped by assimilating it to a more general kind of conversational pretense.

3 | PREVIOUS PROPOSALS

The main aim of this essay is to argue that semantic pluralism provides an appealing way to make sense of assertoric imprecision—one that is compatible with both (TRU) and (SEM). In this sense, I aim to defend the following conditional claim: if one accepts both (TRU) and (SEM) but also seeks to make sense of assertoric imprecision, then one should adopt semantic pluralism.

Quite clearly, someone could accept this conditional claim even if they were strongly inclined to reject at least one of (TRU) or (SEM). Still, it will be worth showing that there are good reasons to accept the antecedent of the conditional: that there are genuine theoretical costs associated with giving up on either (TRU) or (SEM) in an effort to make sense of assertoric imprecision. That shall be the business of the present section, in which I present and critically evaluate two kinds of responses to the problem of assertoric imprecision. The first consists in a rejection of (TRU) and the second a rejection of (SEM). Although these accounts are compatible and could fruitfully be combined, I'll treat them as distinct for the purposes of discussion. While each of these approaches captures something important about the phenomena under consideration, I will argue that they suffer from serious difficulties. This discussion will then set the stage for the development of my own view in Section 4.

3.1 | Relevance theory

Perhaps the most straightforward response to the problem is to simply deny that truth is a norm governing assertion. This style of response drops (TRU): the claim that, as a default, one should only make an assertion provided that one takes the content expressed by that assertion to be true (in the world of evaluation). According to this view, there is nothing particularly surprising or theoretically significant about cases in which speakers assert falsehoods: that is, just how conversation is designed to function. Such a claim is, of course, compatible with the fact that the contents of some assertions are true and that truth is something for which speakers often aim in making their assertions. What the view denies is there are constitutive connection between truth and assertion.

The challenge for any response of this sort is to provide an account of the norms of assertion that make no essential appeal to alethic notions like truth or knowledge. In a series of important papers, Deidre Wilson and Dan Sperber have attempted to address this challenge head-on using the framework of *Relevance Theory*. Relevance theory is a general theory of interpersonal communication. The mechanisms employed by the theory are aimed at modeling a process that applies after a hearer has decoded the syntactic and semantic properties of the speaker's utterance; the principles of Relevance

Theory are supposed to take the hearer from those syntactic and semantic properties to the speaker's intended meaning.⁷

According to Wilson and Sperber, the problem of assertoric imprecision shows that verbal communication is not governed by a norm of truthfulness but instead is governed by a *norm of relevance*. In their framework, relevance is determined by a distinctive kind of cognitive cost–benefit analysis: relevance is a feature that holds for cognitive processes that maximize positive cognitive effects, while minimizing processing costs. In the case of verbal communication, relevance is given by a maximization of the information an utterance conveys while minimizing the cost of compression on the part of hearers. Relevance is thus an essentially comparative notion, varying with the background assumptions, interests, and cognitive capacities of those involved in a conversation; it is also is a matter of degree, something is more or less relevant but not relevant *simpliciter*.⁸

Wilson and Sperber define a property of *optimal relevance* that applies to the interpretation of speakers' utterances:

(OR) An interpretation I of utterance U is optimally relevant for hearer H iff I is relevant enough to be worth H's processing effort, and that content is the most relevant one compatible with the speaker of U's abilities and preferences.

For Wilson and Sperber, linguistic communication is driven by the attempt to maximize optimal relevance; it is optimal relevance that speakers aim for, and which hearers expect. Truthfulness, in this model, is at best an interaction effect: it is among the positive cognitive effects of what is communicated, but truthfulness is not the only positive effect, and it is not the driving force behind linguistic communication in general or the interpretation of assertion in particular. Relevance theory thus seems well suited to address the problem of assertoric imprecision. Imprecise assertions are ways to make maximally relevant contributions to the conversation, even though their contents are not strictly true.

Relevance Theory offers a novel way to develop a non-alethic approach to assertion. I believe that Wilson and Sperber are correct in emphasizing the ways in which the interpretation of speakers' linguistic utterances is sensitive to a background network of intentions, expectations, and preferences. As I shall go on to argue, this is especially important for characterizing the phenomena of assertoric imprecision. However, their rejection of a substantive role for truth in the process of interpretation is misguided. I will raise one objection that applies to any response to the problem of assertoric imprecision that rejects a substantive role for truth in the process of linguistic interpretation, and a second objection that is specific to Relevance Theory.⁹

⁷Relevance theorists distinguish between two such processes: one which allows the hearer to determine the intuitive truth conditions of the speaker's utterance (what they call the *explicatures* of the utterance), and another which allows the hearer to determine the information a speaker intended to convey over and above the truth conditions of the utterance (what they call the *implicatures* of the utterance).

⁸I have followed Sperber and Wilson in defining relevance to be a property of the inputs to cognitive processes. However, this formulation is ambiguous between the claim that relevance is a property of utterances themselves and the claim that relevance is a property of possible interpretations of utterances. These two formulations are clearly not equivalent, as the latter makes relevance hearer-dependent in a way that the former does not. While Sperber and Wilson are not entirely clear of this point, I shall assume that relevance is a property of possible interpretations of utterances.

⁹Another worry that is specific to the relevance theoretic treatment of imprecision is that it has a difficult time explaining the effect of slack regulators. Slack regulators have clear effects on the precision of the contents with which they compositionally interact. But the principles of relevance are supposed to operate after the principles of syntactic and semantic composition have run their course. Whether a version of relevance theory could be extended to provide a treatment of slack regulators in an interesting and very much open theoretical question, but one which I will leave for another occasion.

One general worry for accounts that reject a constitutive connection between truth and assertion is that they fail to do justice to the distinctly epistemic dimensions of assertion.¹⁰ Assertoric uses of language do not merely serve to communicate information, they also serve to provide others with epistemic warrant concerning the information communicated. That is to say, speakers use assertions to put one another in a position to acquire knowledge. There has been much dispute concerning exactly what is required in order for epistemic warrant to transfer across agents in testimony. I assume, however, that the following are at least necessary conditions: (i) hearers must be able to reliably recover the content that a speaker's speech act served to express, and (ii) the content that is expressed by a speaker's speech act must reliably co-vary with the states of the world. Accordingly, in order for epistemic warrant to transfer through testimony there must minimally be reliability in both the mind-to-mind connection and the world-to-mind-to-mind-to-world connection.

Accounts that reject a constitutive connection between truth and assertion, like Relevance Theory, have no trouble meeting condition (i): they can account for the reliability of communication. However, they do not ensure that condition (ii) is met: since assertion is not governed by a norm of truth or knowledge, the content communicated by a speaker's assertion need not co-vary with the states of the world. This threatens to undermine the epistemic significance of testimony, since it would allow hearers to acquire testimonial knowledge only from subset of speakers—namely, from those speakers who care about being truthful or who audience members have acquired specific reason to trust. The trouble is that, as a number of epistemologists have argued (c.f. Burge, 1993, 1997), we acquire knowledge from others' assertions as a *default*. To return to the case with which this paper began: Bea can acquire knowledge of the time of day from Alf's assertions even if she has no special reason to believe that Alf is more concerned to be truthful or a trustworthy source than any other stranger she might have encountered. Those who give up (TRU) have a difficult time making sense of the default epistemic warrant we receive from others' assertions. Although some—including Sperber himself (Sperber et al., 2010)—may be happy to accept this consequence, I take it to be a serious cost.

The second related worry is specific to the pragmatic framework developed by Sperber and Wilson. Relevance Theory is, as I've said, a general theory of interpersonal communication. Wilson and Sperber's treatment of assertoric imprecision focuses on the process by which audience members attempt to recover a speaker's intended meaning. Given this fact, Wilson and Sperber seem to make a prediction: the phenomena at issue should only arise in cases of interpersonal communication. For if the phenomena could arise in cases outside of interpersonal communication, then the solution offered would not be sufficiently general. But the prediction made by Wilson and Sperber is mistaken; the problem at issue can equally well be raised in the case of the contents of thoughts—for instance, in Sam's *believing* that he arrived to the party at 6:00 when he sees that his watch says it is 6:03 or in my *remembering* that there were two hundred students in the lecture when I was told that one hundred ninety-eight students were there. There is, in other words, just as much a problem of *attitudinal imprecision* as there is a problem of assertoric imprecision. But Relevance Theory does not obviously apply to the processes that serve to generate mental content. Indeed, it is not even clear what it would be for a thought content to be generated through a process of optimal relevance. Apart from cases of self-reflection, one does not typically *interpret* the content of one's own thoughts: *one simply has them*.

These two problems are revealing. The first problem suggests that an adequate account of assertoric imprecision should retain the connection between truth and assertion. The second problem suggests that an adequate account of assertoric imprecision should center on features of the semantic

¹⁰A related argument has been developed by Sanford Goldberg; see Goldberg (2001).

contents of speakers' assertions so that the account could be generalized to apply other systems of semantic representation besides speech acts—for example, to the contents of thoughts.

3.2 | Semantic minimalism

A second, and in my view more plausible, response is to retain the claim that assertion is governed by a norm of truth, but to maintain that this norm does not govern the literal meaning of speakers' assertions; instead, the truth norm applies to propositions that speakers convey apart from the literal meaning of their words. This style of response drops (SEM): the claim that, as a default, the content expressed by an assertion of sentence *S* in context *C* is the conventionally determined semantic content of *S* in *C* or its literal meaning. With respect to the problem of assertoric imprecision, the central idea is that this will provide us with the resources to explain how speakers could successfully convey truths in making loose assertions, despite the fact that the literal meaning of the sentences asserted (their semantic content) is false.

Lasersohn (1999) has developed an account along these lines to address the problem of assertoric imprecision. Lasersohn's discussion is distinctive because he attempts to provide a compositional procedure for calculating the propositions pragmatically conveyed—by not literally expressed—by a speaker's assertion. This latter feature is especially important since views which give up on (SEM) have often be criticized for failing to identify clear principles for explaining how speakers and hearers manage to coordinate on determined propositions which are merely pragmatically conveyed but not semantically expressed.

Lasersohn proposes an account in which uses of expressions are associated with both a semantic value and a set of *pragmatic halos*. A pragmatic halo for an expression *E* is a set of denotations that differ from the actual semantic value of *E* in ways that are pragmatically ignorable in the context in which *E* was used. More exactly:

PRAGMATIC HALO: For any expression α , the pragmatic halo of α is a set of objects of the same semantic type as $[[\alpha]]_{w,c}$ that differ from $[[\alpha]]_{w,c}$ in ways that are pragmatically ignorable in *C*.

So, for example, an expression like “Chicago” will receive a region as its semantic value in context, and its pragmatic halo will consist of a set regions that differ in pragmatically insignificant ways in that context from the semantic value of “Chicago.”

It is worth emphasizing two points about pragmatic halos. First, the set of pragmatic halos associated with an expression will be ordered, either totally or partially, with the semantic value of that expression forming a center for the ordering. This is to say that objects comprising a pragmatic halo are ranked according to the degree to which they differ from the semantic value of the expression in question. This point is important because Lasersohn maintains that slack regulators—or expressions which serve to modify the precision of the content of other expressions with which they compositionally interact—serve to affect the ordering associated with a pragmatic halo. Expressions like “precisely” and “exactly” are held to narrow a halo to those closer to the center, while expressions like “approximately” or “around” are said to expand a halo to incorporate values further away from the center.¹¹

¹¹In an important recent discussion, Carter (2021) argues that slack regulators operate in a compositionally more complex manner than Lasersohn suggests. These complications are important, but they do not substantially impact the discussion that follows.

Second, the pragmatic halo associated with an expression will vary with conversational context. Although Lasersohn says very little about the determinates of what is pragmatically ignorable in a context, it is presumably a function of the presuppositions, intentions, and goals of the conversational participants.¹² As the presuppositions, intentions, and goals of conversational participants shift, so too will the pragmatic halo associated with expressions in context. This is important because it allows Lasersohn to account for the way that the imprecision of expressions varies across conversational situations.

The pragmatic halos associated with expressions will have a systematic impact on what is conveyed by asserting sentences containing those expressions. The semantic content of a sentence in a context will be given in the usual way by composing the semantic values of its parts in context and their manner of combination. Likewise, the pragmatic halo of a sentence in context will be given by composing the pragmatic halos of its parts in context and their manner of combination. In asserting a sentence, a speaker both expresses the semantic content of that sentence and in addition conveys a set of propositions given by the pragmatic halo of the sentence. The beauty of this approach is that it provides a compositional mechanism whereby a speaker may assert multiple propositions in uttering a single sentence. This allows Lasersohn to retain the claim that truth is the norm of assertion. In particular, it leads him to maintain that one should only assert a sentence if the pragmatic halo associated with that sentence in context contains at least one true proposition. So understood, the norm of assertion can be satisfied even when a speaker asserts a sentence that semantically expresses a falsehood; for the norm of assertion merely requires the truth of one of the propositions in the pragmatic halo of a speaker's utterance, rather than the truth of the semantic content of that utterance.

Lasersohn's takes this framework to straightforwardly address the problem of assertoric imprecision. Consider an assertion of (3):

- (1) "8,000 people attended the rally today."

The semantic content of (3) is a function which returns the value TRUE at a possible world just in case exactly 8000 people attended the rally in question at that world. If, as a matter of actual fact, there were only 7988 people at the rally, then an assertion of (3) is literally false. But (3) may still be assertable in that context, so long as the difference between the number of people present at the rally is ignorable given the practical interests of the conversational participants. For in such a context, 7988 will be a member of the pragmatic halo associated with "8000," and it is indeed true that 7988 attended the rally. This is supposed to explain why speakers can appropriately assert (3) in spite of the fact that it has a semantic content that is literally false.

Lasersohn's account is elegant, and I believe that it highlights a fundamental feature of the phenomena of assertoric imprecision: namely, that speakers put into play multiple propositions in asserting a single sentence. Nonetheless, Lasersohn's account suffers from a number of serious problems. These problems relate to tacit assumptions that Lasersohn makes concerning the facts in virtue of which linguistic expressions have the conventionally determined semantic contents that they do. In other terms, Lasersohn's account requires highly implausible claims concerning the *meta-semantic* or *foundational* facts that determine the semantic content of an expression in context. I will focus on two related instances of this problem.

First, the account that Lasersohn provides of the truth conditions determined by the semantic content of assertions is extremely demanding. According to Lasersohn, the vast majority of speakers'

¹²See Lauer (2012) for an attempt to develop such an account on Lasersohn's behalf.

assertions semantically express propositions that are false. To provide just one concrete example, his proposed semantics for an assertion of a sentence like (1):

(1) The bell rang 6:00.

is that the expression “6:00” has a single instant of time as its semantic value. But, as Lasersohn notes, it is natural to suppose that events like ringing take more than a single instant to perform—they occur across intervals of time. For this reason, assertions of sentences like (1) are never used by ordinary speakers to semantically express truths in the actual world. Although this specific example could be addressed by providing a different semantics for “ring,” it highlights the quasi-skeptical feature of Lasersohn's account of semantic content. This quasi-skeptical feature of Lasersohn's view has a cost—a cost serious enough, in my opinion, to warrant rejection of the view. Call this the *problem of too-high-standards*.

Lasersohn acknowledges this problem for his account, but he is undeterred. This is because he maintains that our intuitions about acceptability track what is pragmatically conveyed, rather than what is semantically expressed.¹³ While a sentence like (1) may only ever be used to assert a literally false proposition, the account of pragmatic halos allows us to see how that sentence could regularly be used to convey true propositions. Although this way of understanding the character of ordinary speakers' truth-value intuitions raises a host of worries of its own, I want to set the entire issue aside. For, in my view, the real problem lies elsewhere. The problem is that Lasersohn's view is at odds with a familiar meta-semantic constraint according to which we should assign semantic values to expressions in context in a way that is *weighted toward the truth*. For instance, Donald Davidson (1973) and Lewis (1974) require that semantic values be assigned *charitably*, so as to maximize the truth (and rationality) of speakers' utterances. While subsequent discussions have refined and revised these proposals, most theorists have accepted some version of the principle of charity. However, if meta-semantic constraints such as these are at all plausible, then there is a serious problem with Lasersohn's story: the principle of charity leaves little room for views according to which the semantic contents of speakers' utterances are rarely, if ever, true. The trouble with Lasersohn's account is not simply that it flies in the face of our ordinary intuitions about standards for truth but that it is inconsistent with plausible principles of meta-semantics.¹⁴

Lasersohn has another meta-semantic problem: one that pertains to the relationship between linguistic meaning and *use*. It is long been taken to be something of a truism that, in one way or another, the facts concerning the semantic contents of expressions in context are at least partially grounded in facts about the linguistic practices of those producing and responding to those expressions in those contexts. It is difficult to see how Lasersohn can make good on this connection between semantics and use. As we have seen, Lasersohn maintains that the norms of assertion do not govern the semantic contents of speakers' assertions; rather, the norms of assertion govern what is pragmatically conveyed by speakers' assertions. Moreover, semantic contents are neither what speakers are primarily intending to communicate with their assertions nor what audience members need to arrive at in order for communication to be successful. Instead, communicative success is measured by the degree to which the members of the conversation coordinate on pragmatic halos associated with a speaker's utterance.

To be clear: the problem here is not that Lasersohn has failed to say exactly how semantic values relate to linguistic practice or how meaning depends on use. The connection between meaning and

¹³See Lasersohn (1999), p. 544 for discussion.

¹⁴This meta-semantic problem and the problem of use below both apply to more recent treatments of assertoric imprecision that deny (SEM), for example, the accounts suggested by Hoek (2018) and Carter (2021).

use is a notoriously complex and controversial affair, and for many theoretical purposes it is appropriate to abstract away from the features of linguistic practice in virtue of which expressions come to be the “intended interpretations” of linguistic expressions in the universe of the semantic model. The problem is not that Lasersohn omits crucial details relating semantics and use, but rather that he is committed to semantic details that makes semantics and use *unrelatable*.¹⁵ At the very least, Lasersohn owes us some story about what it is about our practices of using language in virtue of which some semantic theory *T* is the one that we use and some alternative semantic theory *T** is not.¹⁶ I'll call this the *problem of use*.

Both Lasersohn's appeal to pragmatic halos and Wilson and Sperber's appeal to Relevance Theory are intended to resolve the problems raised by cases of assertoric imprecision by appealing to pragmatic mechanisms that operate *after* the computation of literal meaning has run its course. I suggest that this is the source of the problems we have encountered for both views. In what follows, I suggest a different sort of account: one that centers on mechanisms that operate *pre-semantically* or prior to the determination of literal meaning. More specifically, I will argue that the problem of assertoric imprecision motivates a version of *semantic pluralism*.

4 | SEMANTIC PLURALISM

The views I have discussed thus far have all tacitly assumed the doctrine of *semantic uniqueness*. Semantic uniqueness is the thesis that background linguistic conventions (or whatever features fix the “actual language relation”) associate a single semantic value with each expression of the language relative to a context of use. I claim that rather than motivating a rejection of (TRU) or (SEM), the phenomena of assertoric imprecision motivates a rejection of semantic uniqueness. In this section, I will develop an alternative framework that I call *semantic pluralism*. After developing this framework in detail, I illustrate how it has the resources to capture the distinctive truth conditions of loose assertions.

The animating idea of semantic pluralism is the claim that there are degrees of freedom in what is required in order for a speaker to correctly apply specific expressions. Speakers and their audiences are responsive to this flexibility concerning the interpretation of the terms in question and exploit it in their conversational exchanges. When speakers assert sentences containing terms that allow for this freedom, they semantically express multiple propositions.

¹⁵In might be maintained that by building on the discussions in Soames (2002, 2009) Lasersohn could respond to this challenge by treating semantic contents as *least common denominators* of the information conveyed by uses of linguistic expressions—that is, as propositions (or components thereof) that are invariably included in the set of pragmatic halos associated with sentences in context. Some of what Lasersohn says suggests that this is precisely how he is thinking of semantic content. Unfortunately, however, the least common denominator approach to semantics is highly implausible. For one thing, the account flounders on context-sensitive expressions: there are no plausible entities in the world that serve as least common denominators for every use of expressions such as “you” or “that.” For another, when coupled with Lasersohn's other commitments this approach has resulted that the semantic content of a sentence in a context is the one proposition conveyed by the use of that sentence that speakers never truthfully assert or are primarily intending to communicate. See Stanley (2005) and Stanley and King (2005) for further discussion.

¹⁶The problem is not that it is difficult to explain how the facts about the use of language among a group of agents could determine the pragmatic halos associated with an expression in context or even to explain how such pragmatic halos could come to be encoded in the form of a sentential operator or subsentential semantic modifier. The problem is to explain what it is about the linguistic practices of such agents in virtue of which their *unembedded* and *unmodified* linguistic expressions in context have the literal, truth-conditional contents that they do rather than some alternative set of truth-conditional contents.

Semantic pluralism is not a novel doctrine—it has been suggested by number of contemporary authors, in a general form by Hawthorne and Dorr (2014) and in more specific cases by other authors.¹⁷ Indeed, a version of semantic pluralism was initially suggested by David Lewis in a neglected passage of *Convention* (1969):

I think we should conclude that a convention of truthfulness in a single possible language is a limiting case—never reached—of something else: a convention of truthfulness in whatever language we choose of a tight cluster of very similar possible languages....The sort of convention I have in mind is this: almost everyone, almost always, is truthful in at least some languages in the cluster: but not necessarily the same ones for everyone or for one person at different times. (1969, p. 201).

I believe that Lewis' suggestion here can be exploited to resolve the problem of assertoric imprecision. Unfortunately, Lewis never systematically connected his discussions of the dynamics of conversation with his discussions of conventional meaning. In particular, Lewis did not seem to appreciate the degree to which the “tight cluster of possible languages” established by linguistic conventions is affected by highly dynamic features of the context of utterance. This dynamic interaction between literal meaning and context is crucial for understanding cases of assertoric imprecision.

We can gain theoretical traction here by recognizing that members of a conversation have mutual attitudes concerning what the world is like but also about how they will interpret one another's uses of linguistic expressions. These mutual attitudes vary across conversations, depending on the background assumptions, interests, and questions of conversational participants. And these mutual attitudes also undergo change within conversations as participants produce and respond to one another's speech acts and gain mutual information about the world through other means.

In terms of the model of common ground discussed in the first section, I am proposing that our model of the common ground should be enriched to include a set I of *interpretation functions*—functions which each provide a mapping from the simple and complex expressions of the language to intensions of the appropriate semantic type for each expression in the language.¹⁸ So rather than characterizing the common ground in the usual Stalnakerian fashion in terms of a context set or set of possible worlds, the common ground will here be characterized in terms of a set C of ordered pairs $\langle w, i \rangle$ consisting of a possible world and an interpretation function.

Moreover, just as some worlds in the common ground are treated by agents as live or open possibilities, so, too, some interpretation functions in the common ground are treated as *admissible* or as being *live candidates for interpretation*. Some such interpretation functions will be *globally admissible*—that is, admissible in all (or nearly all) conversational contexts.¹⁹ These globally admissible interpretation functions respect what Kit Fine (1975) has called “penumbral connections” holding between the expressions of the language. In contrast, other interpretation functions will only be *locally admissible*

¹⁷See von Fintel and Gillies (2011) for the cases of epistemic modals, King (2014) for the case demonstratives, and indeed context-sensitive terms generally, and Tappenden (1993), Soames (1999), Shapiro (2006), and especially Barker (2002) and Raffman (2014) for the case of vague expressions.

¹⁸I assume that the interpretation functions in I will be characterized recursively in a familiar manner so that, for example, each interpretation function will assign an individual at a world to names, n -ary sequences of individuals at a world to n -place predicates, a truth value at a world to sentences, etc. But this assumption is not required for any of the claims I make in what follows.

¹⁹I take the difference between global admissibility and local admissibility to be a difference of degree rather than a difference in kind. In particular, there may well contexts of conversation or of inquiry—say, in the logic seminar room or in exploring the contents of Borges' imagination—in which global constraints on admissibility are relaxed or completely revised.

or admissible in light of an agent's or group of agents' specific expectations about how the expressions of a language are to be interpreted in some, but not all, conversational contexts. In this sense, locally admissible interpretations will vary depending on agents' expectations about prior uses of language, about those they are interacting with, the topic of conversation, the questions under discussion, and other features of the conversational context.

My claim is that agents' interpretative attitudes toward their language are of *semantic* significance. Semantic values are associated with speakers' use of simple and complex expressions in accordance with what is admissible given the background state of the common ground.

For some linguistic expressions, the background state of the common ground may well determine a uniquely admissible semantic value for uses of those expressions. For example, given the conversational participants' mutual attitudes, it may be that the man Val Kilmer is the only admissible interpretation of uses of the name "Val Kilmer." But what is distinctive about semantic pluralism is the claim that the common ground need not—and often will not—determine a unique semantic value for each expression of the language. In particular, the theory allows that certain expressions may be associated with *set* of semantic values. When speakers assert sentences containing such expressions, they will semantically express multiple propositions (or a set of sets of possible worlds) in uttering a single sentence. In such cases, the relationship between linguistic expressions in context (e.g., uses of expressions) and semantic values is one-to-many rather than one-to-one.

What does semantic pluralism suggest about athletic norms on assertion, like those at work in (TRU)? While semantic pluralism is compatible with a number of different attitudes toward the norms on assertions, I suggest that it leads naturally to the following set of claims. The norms on assertion govern the total set of propositions expressed by speakers' assertion. What the literal truth of a speaker's assertion requires is that at least *one* of the propositions that speaker semantically expresses is true. More carefully, for an utterance of a sentence s in context c to be true, there must be some admissible $\langle w, i \rangle$ pair in C such w is the actual world and $i(s)$ is true with respect to w . In contrast, a speaker's assertion of a sentence is literally false just in case *none* of the propositions that speaker semantically expresses are true. In other words, the relevant norm on assertion requires that speakers are committed to the actual world (of evaluation) being in the union of the set of set of worlds that their assertions serve to semantically express. In this sense, the alethic evaluation of speakers' assertions can be said to be *charitable* or *weighted toward the truth*.²⁰

Consider an example. Suppose that the common ground of a conversation allows for three admissible interpretations of uses of the expression "6:00"—say intervals t_1-t_2 , t_2-t_3 and t_3-t_4 . Further suppose that all other expressions of the language have uniquely admissible interpretations. Now consider an assertion of (1) with respect to that state of the common ground:

- (1) The bell rang at 6:00 p.m.

In asserting (1), the speaker semantically expresses multiple propositions: the proposition that the bell rang to the party at t_1-t_2 , the proposition that the bell rang at the party at t_2-t_3 , and the proposition

²⁰Notice that the conditions for literal truth and falsity are not symmetrical. The literal truth of assertion requires the truth of *a* proposition semantically expressed, whereas the literal falsity of an assertion requires the falsity of *every* proposition semantically expressed. In the terminology of Fine (1975), assertoric truth is *not* identified with SUPERTRUTH, while assertoric falsity *is* identified with SUPERFALSE. An alternative to the present account would be to give a more symmetrical treatment by embracing a paraconsistent logic. Indeed, semantic frameworks of this sort have been developed under the label of subvaluationism; See Hyde (2007), and Ripley (2013) for discussion.

that the bell rang at t_3-t_4 . The conversational effect of asserting this cloud of propositions will be akin to asserting an explicitly disjunctive proposition which states the bell rang at one of the admissible interpretations of '6:00.'²¹ Likewise, if audience members accept the content of the speaker's assertion, they will eliminate possibilities according to which the bell rang at times outside the range of admissible interpretations—in the case, at any interval apart from t_1-t_2 , t_2-t_3 or t_3-t_4 .

This example is highly simplified. It involves a case in which there are a small number of interpretative possibilities involving just a single expression, and it is informally described. But the simplistic example helps to illustrate the way in which the model I am suggesting offers an alternative conception of what is required to satisfy the truth conditions given by the literal meaning of a speaker's words.

In order for the assertion of a sentence to be literally true, the actual world (or world of evaluation) must be a member of the union of the set of sets of possible worlds semantically associated with that sentence in the context in which the assertion was made. In the toy example of an assertion of (1) just given, the speaker's assertion is literally true just in case the bell rang at t_1-t_2 , or t_2-t_3 , or t_3-t_4 —the bell arriving at any one of those times in the actual world is sufficient to make the speaker's assertion literally true. In contrast, the literal falsity of an assertion requires that the actual world of evaluation is not among the union of the set of set of possible worlds semantically expressed by a speaker's assertion.

As this simplistic example illustrates, assertoric contents may be *literally true* without *precisely* corresponding to the facts. We can (roughly) measure the imprecision of an expression in terms of the variation present in its admissible interpretations in a context. The occurrence of an expression is precise, in this sense, if that expression has a uniquely admissible interpretation, given the state of the common ground in which it was used. But speakers can make assertions that are literally true but which also contain imprecise occurrences of expressions. Given the context in which it was asserted, the semantic content of (1) is imprecise: it consisted of a set of admissible propositions. But the speaker nonetheless was being literally truthful in asserting (1), because one of the propositions semantically expressed in the context of utterance was indeed true of the circumstances being evaluated.

Does this proposal require any serious revisions to (TRU)? It does not. In my view, it is a mistake to construe norms on assertions—like (TRU)—as requiring speakers to make assertions that are both true and fully precise. If a version of semantic pluralism is accepted, (TRU) could be more aptly stated as the thesis that *as a default, one should only make an assertion provided that one takes at least one proposition expressed by that assertion to be true*. In the special case in which there is a single proposition expressed by a speaker's assertion, this norm will require that the speaker takes this single proposition to be true. But in cases in which there are multiple propositions expressed by a speaker's assertion, the norm will be satisfied so long as at least one proposition asserted by the speaker corresponds to the way things are in the world of evaluation.

One reason why this way of construing the norm of assertion is not revisionary is that it preserves one of the core reasons why we care about the truth of assertions in the first place. We care about the

²¹There is a useful if potentially misleading analogy to be made between the effect of asserting an unembedded disjunctive sentence and the effect of making an imprecise assertion on the account I have developed. The analogy is useful because it provides a model for the truth conditions I am proposing for imprecision assertions: just as the truth of one of the disjuncts is enough to make an assertion of an unembedded disjunctive sentence true, so, too, the truth of one of the propositions expressed is enough to make an assertion of an imprecisely used sentence true. However, the analogy is potentially misleading because there is often (though not always) higher-order imprecision concerning the set of propositions associated with a speaker's imprecise assertion. Quite commonly, there will be borderline cases concerning where the range of admissible semantic values governing the use of an expression begins and ends. In this respect, the truth conditions for imprecise assertions might more aptly be analogized to the truth conditions for an utterance of an *open-ended* unembedded disjunctive sentence.

truth of a speaker's assertion because, among other things, we want to be confident that updating on a speaker's assertion will not lead us to eliminate the actual world (of evaluation) from active consideration in our practical and theoretical pursuits.

The present framework allows us to vindicate this concern. If a speaker's assertion of a sentence is true, as construed by the lights of semantic pluralism, then the result of updating on the union of the set of propositions expressed by that speaker's assertion will not result in the actual world being excluded from active consideration in the conversational context. Indeed, updating on the speaker's assertion will serve to narrow the live possibilities in the common ground to those that more closely approximate the way things actually are in the circumstances under consideration. After all, an asserted sentence s updates a context C by eliminating all pairs $\langle w, i \rangle$ from C in which w is not in $i(s)$ (cf. Stalnaker, 1978). And, from what we have said earlier, an assertion of a sentence s is true in a context c just in case, for some i in I at c , the actual world is in $i(s)$. Hence, so long as the context contains an admissible pair consisting of an interpretation function that maps s onto a true proposition and the actual world, that pair will not be eliminated by the update.²² In this way, we can preserve what's important about alethic norms on assertion while also not demanding more semantic exactness than is required by the goals of the conversational participants.

The present framework also provides a way to preserve (SEM) in spite of the facts about assertoric imprecision. There may well be theoretical purposes—say, in characterizing recursive principles of phonological, syntactic or semantic composition—for which it is a useful idealization to treat literal meaning as being invariant with respect to a group of agents' interpretative attitudes concerning uses of language. But if taken too seriously, this idealization can distort our understanding of linguistic practice. Our practice of using linguistic expressions imprecisely is a case in point. The admissible interpretations of the linguistic expression “Chicago” in a casual conversation at an airport is significantly different from what is an admissible interpretation of that linguistic expression among those carrying out a city census. Similarly, absent special circumstances, there will be clear differences between the admissible interpretations of ‘7997’ in an assertion of (6) and the admissible interpretations of “8000” in an assertion (3):

(6) 7,997 people attended the rally today. (3) 8,000 people attended the rally today.

These difference in the interpretation of expressions in context are conventionalized among typical speakers of English: although arbitrary, they perpetuate over time because of the interlocking and mutually reinforcing attitudes of language users.²³ By taking these facts about our practices of linguistic interpretation to have semantic import, we can give a natural treatment of assertoric imprecision without thereby severing the connection between literal meaning and use.

5 | SEMANTIC PLURALISM AND ASSERTORIC IMPRECISION

I have shown how the framework of semantic pluralism offers the resources to address the intuitive problem of assertoric imprecision. In this section, I return to the two theoretical problems concerning assertoric imprecision noted in Section 2—specifically, the data involving abominable conjunctions

²²Of course, the actual world might come to be eliminated from active consideration even by updating on a true assertion if audience members happen to misinterpret a speaker's assertion and invoke an interpretation function that maps the sentence asserted onto a false proposition.

²³In Armstrong (2016), I offer a more concrete characterization of linguistic conventions that bears out this claim.

and retractions. I argue that these data does not provide compelling reason to reject the framework I have developed.

5.1 | Abominable conjunctions

Recall that one of the desiderata of an account of assertoric imprecision was to explain what is wrong with asserting sentences like (4) and (5):

(4) #“Although the bell rang at 6:00, it wasn't until 6:03 that it rang.”(5) #“I am going to a conference in Chicago this weekend; it is in Evanston, just outside of Chicago.”

These sentences are clearly infelicitous. Indeed, it is difficult to hear any true readings of assertions of these sentences. And recall, also, that assertions of (4) and (5) stand in marked contrast to (4') and (5'):

(4') “At approximately 6:00 the bell rang. Indeed, it rang at 6:03.”(5') “I am going to a conference around Chicago this weekend; it is in Evanston, just outside of Chicago.”

Assertions of these sentences seem totally felicitous. Indeed, it is not difficult to hear true readings of assertions of these sentences.

It may be alleged that my view makes incorrect predictions about these cases. After all, I have claimed that a speaker may assert a sentence such as “The bell rang at 6:00” and express a true semantic content even if, in actual fact, the bell rang at 6:03. This would seem to suggest that assertions of (4) and (5) should be perfectly felicitous and indeed true, just as we find assertions of (4') and (5').

I submit that assertions of sentences like (4) and (5) are not true, and not predicted to be true by the account I've developed. There are no true readings of these sentences, because no admissible interpretation function is such that it allows (4) or (5) to be true. This is because assertions of sentences like (4) and (5) violate natural penumbral connections that govern the interpretations of terms. In the cases at hand, these penumbral connections ensure that there is no pair of admissible worlds and interpretation functions in the context of utterance which would allow a single event to be truly described as simultaneously taking place at two separate times (or intervals of time) or for a single location to be truly described as having two distinct spatial configurations at once. Principles of this sort would explain why—absent extremely special conversational contexts—it is never true for a speaker to assert “it's 3:00 and it's not 3:00.” And this is so even in contexts in which there are pairs of propositions *P* and *Q* such that the first conjunct semantically expresses *P* (perhaps, amongst other propositions) and the second conjunct semantically expresses *Q* (perhaps, again, amongst other propositions), and the conjunction of *P*&*Q* is not a contradiction.²⁴

This explanation can be extended to sentences like (4), given a plausible constraint on admissibility according to which distinct expression types in the language are expected to be associated with distinct

²⁴There is a puzzle that remains. This is the question of why we cannot explicitly assert pairs of sentences that serve to express propositions which are both consistent with one another and compatible with what's admissible given the common ground. I suspect that we can make progress on this puzzle by reflecting on the ways in which the act of explicitly asserting a sentence in a context can serve to change what is admissible in the subsequent context. However, a proper treatment of this interesting issue will have to wait for another occasion.

semantic interpretations.²⁵ If each admissible interpretation function in the common ground is such that “3:00” and “3:03” are associated with a different set of semantic values, then it is not surprising that a sentence like (4) cannot be truly asserted: a speaker would be asserting that there was a single event of a bell ringing that took place at a range of times associated with “3:00” and also at a distinct range of times associated with “3:03.” But there is no admissible world and admissible interpretation function in the common ground which would allow a single event to take place at multiple times or, in the case of an assertion of (5), for a single location to have multiple distinct spatial configurations at the same time. This would provide a straightforward explanation for the “abominable” character of these assertions: we judge assertions like (4) and (5) to be unacceptable and, indeed, contradictory, because there are no pair of factual and linguistic possibilities available in the context which would serve to make them true.

Notice that an assertion of a sentence like (4′) does not violate conditions on admissibility in this way. Part of the point of including slack regulators like “approximately” or “around” is to explicitly signal that the relevant events did not happen over the very same intervals of time and that the relevant spatial locations do not perfectly overlap. Thus, while penumbral principles may prevent assertions of (4) or (5) from having any true readings, those principles do not prevent assertions of (4′) or (5′) from having true readings. This is as it should be, and the present account can capture it.

Some may be unconvinced by these considerations, holding that the appeal to penumbral connections is ad hoc or raises as many problems as it resolves. Let us therefore set this line of response aside, at least for the sake of the present argument. Suppose that there were semantically acceptable, non-contradictory, readings of assertions of (4) and (5). What that provide a decisive reason to reject the present account? It would not. This is because the unacceptability of assertions of (4) and (5) can also be explained *pragmatically* through an appeal to general norms on language use. There are plenty of assertions that are perfectly consistent—and perhaps even true—that are nonetheless not assertable. As an example, consider an assertion of (7):

(5) #The tour bus will stop in Wyoming or Utah, but it will not stop in Wyoming.

An assertion of (7) is infelicitous. Indeed, this assertion is infelicitous even if we stipulate that a truth-functional semantics for “or” is common knowledge among the conversational participants. In this case, the semantic interpretation of (7) is perfectly consistent—and may well be true—and yet (7) is not assertable. This is presumably because there is a constraint on disjunction to the effect that one should not assert a disjunction and in the same breath remove one of its disjuncts. A related constraint would serve to explain the infelicity of (4) and (5), given the framework I’ve developed. In asserting a sentence like (4) the speaker would be expressing a range of propositions and in the very same speech act removing one of the propositions from within that range. A speaker may be able to do so consistently, as for example in (4′) and (5′), but without special signaling we should expect such assertions to be regarded as pragmatically odd or inappropriate.

The more general point is this: while truth (or knowledge) may be the only constitutive norm governing assertion, there are typically other (non-constitutive) norms at play in conversational exchanges

²⁵Linguists have often called a constraint of this kind *mutual exclusivity* or *the principle of contrast*. There is good empirical evidence—reviewed in Markman (1989) and Clark (2017)—that from a very early age language users robustly expect mutual exclusivity to govern the interpretation of uses of words. For present purposes, it thus does not matter much whether principles of this sort follow from global conditions on admissibility or are instead due to widespread local conditions on admissibility. All that matters for the purposes of this reply is the lack of admissible world interpretation function pairs in the context according to which “3:00” and “3:03” are associated with completely overlapping semantic contents.

involving assertions (Grice, 1967). In addition to aiming to speak truly, speakers also generally aim to be as relevant, orderly, and as precise as possible with their assertion given the goals of a conversation. Assertions may be deemed infelicitous for failing to satisfy these norms, even if they are literally true.

This general point applies not just to single utterances but to separate utterances as well. There is something decisively marked about a situation in which a speaker asserts “The bell rang at 3:00,” pauses a few moments, and continues with, “When the bell rang at 3:03, I was surprised because...”. This is so even if 3:03 was clearly within the set of admissible interpretations associated with “3:00” in the context of the initial assertion.

Even setting aside a direct semantic explanation of this oddness, the speaker's two assertions violate reasonable pragmatic norms on conversation. Just as speakers are expected not to change the topic of conversation or questions under discussion without any observable rationale, so too they are expected not to change the operant standards of precision in a wholly capricious manner.

Standards of precision do change both across and within conversational contexts, as I have emphasized throughout my discussion. However, these changes in the standards of precision are expected to be *principled*: to change on the basis of reasonable shifts in the psychological states of the conversational participants, including shifts that are induced by acts of interpersonal communication or by the unfolding circumstances in the world that are manifest to the conversational participants. But a situation in which a speaker first describes an event as taking place at “3:00” and then describes that very same event as taking place at “3:03” would—without any rationale provided by the speaker or by the changing circumstances in the world—involve a capricious change in the operant standards of precision. This is because, in most conversational settings, uses of “3:03” are associated with a much more constrained set of admissible interpretations than uses of “3:00,” and, in this sense, involve more demanding standards of precision. If the speaker was in a position to use these more demanding standards of precision and deemed them relevant to the goals of the conversation, there would be no clear rationale for why that speaker would first use the less demanding “3:00” only to immediately shift to using “3:03.” The difficulty of being able to identify a natural rationale along these lines provides a purely pragmatic explanation of the oddness of these assertions.

Of course, this is not itself to explain why it is that “3:00” is standardly associated with larger range of admissible interpretations than “3:03” or why “Chicago” typically allows for a wider range of admissible interpretations than “Evanston.” In this section and in the last, I have presupposed that there are these differences associated with our uses of linguistic expression in context and I have presupposed that these differences are due to the psychological profile of agents and their practices of social coordination. However, I have not attempted to isolate the specific psychological or social facts which serve to ground the underlying differences in interpretative degrees of freedom. Although providing a foundational explanation of this kind would enrich the framework I have developed in this paper, it is not itself required to undermine the threat that the data involving abominable conjunctions generate for the claim imprecise assertions can nonetheless be literally true.

5.2 | Retraction data

Let me close this section by address the reaction data. Recall that this data turns on the fact that speakers often concede that they spoke falsely when challenged in cases of assertoric imprecision. The point is illustrated in the following dialogue:

Alf: 8,000 people attended the rally today. Bea: That's not true: there were only 7,998 people at the rally today. Alf: OK, fine, you're right: there were only 7,988 people at the rally today.

Let us assume that prior to Alf's initial assertion, 7908 was among the admissible interpretations of "8000." So according to the view I have proposed, Alf's assertion was literally true. The data involving retractions raise two kinds of challenges for my view: first, to explain why Bea would reject Alf's assertion if the content of his assertion was literally true; second, to explain why Alf would agree that he spoke falsely and retract his assertion in response to Bea's challenge if the content of his assertion was literally true. I will address each of these challenges in turn.

The framework I have developed allows that an assertion may be literally true while nonetheless failing to be fully precise. The distinction between literal truth and full precision is important for understanding the retraction data at hand. I suggest that conversational participants sometimes reject a speaker's assertion on the grounds that the content of the speaker's assertion failed to be adequately precise. Indeed, I think that conversational participants sometimes judge the content of a speaker's assertion to be false because they took the standards of admissibility to be different from the standards of admissibility that guided the speaker's assertion: what the audience took to be admissible was more demanding, and, by those more demanding standards of admissibility, the speaker's assertions were indeed literally false. We should understand Bea's rejection of Alf's assertion in these terms: Bea is rejecting Alf's assertion, and judging it to be false, on the grounds that Alf's assertion failed to true given what she took to be admissible.²⁶

Although her response may be pedantic, Bea is perfectly within her rights to issue such a challenge. As I pointed out in discussing abominable conjunctions, speakers' assertions are typically governed by a variety of norms (even if literal truth is the constitutive norm of assertion); among these norms is that speakers' assertions are expected to be as precise as possible given the goals of the conversation. An audience member, like Bea, can challenge a speaker's assertion if they believe the speaker was not adequately precise given the goals of the conversation. There thus may be good reasons for Bea to reject Alf's assertion, even if what Alf semantically expressed with his assertion is true. In addition, this helps explain why Alf would retract his assertion in response to Bea's challenge; for her assertion signals that Alf did not meet the level of precision that Bea expected. Since Alf's assertion was imprecise, Bea's assertion gives him a reason to retract.

Speakers are not, however, compelled to retract: they can rebuff the challenge and stick to their original assertions (c.f. von Fintel & Gillies, 2010). Rather than retracting, Alf could instead have responded as follows:

Alf: 8,000 people attended the rally today. Bea: That's not true: there were only 7,998 people at the rally today. Alf: Oh, come on, that's close enough: 8,000 people were at the rally.

Some may find Alf's response here less acceptable than others, but it is response that ordinary speakers can and do avail themselves of on occasion. This fact highlights the way that speakers have some

²⁶A full discussion of indirect speech reports—including an account of the truth conditions of indirect speech reports—is beyond the scope of the present paper. However, I assume that each member of the set of admissible interpretations of an assertion can be individually targeted in an indirect speech report; see Hawthorne and Dorr (2014) and Cappelen and Lepore (2015) for relevant discussion of this issue.

individual discretion in how precisely they deem it appropriate to describe a situation.²⁷ The conventions of language and the background state of the common ground can allow for one speaker to categorize a situation in one way and for another speaker to categorize that same situation in a different way. If neither of these speakers is willing to revise their judgments, they will be at an interpretative impasse. But speakers can find themselves at an interpretative impasse without either of them being semantically incompetent or having made a factual error.

Still, I do think there is a real tendency of speakers to retract their imprecise assertions when challenged. This tendency has another source, one that falls out of the dynamics of assertion that the present model highlights. The contents of assertions depend on the prior state of a common ground; but the contents of assertions also serve to change the posterior state of the common ground. If accepted by the other members of the conversation, assertion updates the common ground by excluding factual and linguistic possibilities inconsistent with the content of that assertion. In the case at hand, Bea's assertion was intended to eliminate 7998 as an admissible interpretation of '8000.' Given that Alf did not reject Bea's assertion, we can assume that the common ground was updated accordingly. What follows is that Alf can no longer truly make his initial assertion—given the shifted standards of admissibility, asserting that “8,000 people at the rally today” would be literally false. This does not show that Alf's initial assertion was false. But it does help make clear why Alf might be tempted to retract his assertion.

6 | CONCLUSION

I submit that semantic pluralism is an attractive framework, one that has clear advantages in its treatment of assertoric imprecision in comparison with previous approaches. I have not attempted to address the full complexities of imprecision in natural languages. Rather, I have used simple cases to illustrate the way in which semantic pluralism offers the resources to capture the truth-conditional content of imprecise assertions while retaining a constitutive connection between assertion and truth and between assertoric content and semantic content.

ORCID

Josh Armstrong  <https://orcid.org/0000-0002-0769-4485>

REFERENCES

- Armstrong, J. (2016). The problem of lexical innovation. *Linguistics and Philosophy*, 39(2), 87–118.
- Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, 25, 1–36.
- Burge, T. (1993). Content preservation. *The Philosophical Review*, 102, 457–488.
- Burge, T. (1997). Interlocution, perception, and memory. *Philosophical Studies*, 86, 21–47.
- Cappelen, H., & Lepore, E. (2015). *Liberating content*. Oxford University Press.
- Carston, R. (2002). *Thoughts and utterances: The pragmatics of explicit communication*. Blackwell.
- Carter, S. (2021). The dynamic of loose talk. *Nous*, 55(1), 171–198.
- Clark, E. (2017). Lexical acquisition and the structure of the mental lexicon. In *Oxford Research Encyclopedia of Linguistics*. Oxford University Press.
- Davidson, D. (1973). Radical interpretation. *Dialectica*, 27, 314–328.
- DeRose, K. (1995). Solving the skeptical problem. *Philosophical Review*, 104(1), 1–52.

²⁷This has been noted in many discussions of borderline cases in the literature on semantic vagueness, for example, by Raffman (2014) and Shapiro (2006). Competent speaker can classify an individual as being “bald” in one context but not in another and they can revise their previous classification in the context of a forced-march sorites series.

- Fine, K. (1975). Vagueness, truth and logic. *Synthese*, 30, 265–300.
- Goldberg, S. (2001). Putting the norm of assertion to work: The case of testimony. In J. Brown & J. Cappelen (Eds.), *Assertion* (pp. 175–195). Oxford University Press.
- Grice, H. P. (1957). Meaning. *Philosophical Review*, 66, 377–388.
- Grice, H. P. (1967). Logic and conversation. In P. Grice (Ed.), *Studies in the Way of Words* (pp. 41–58). Harvard University Press.
- Hawthorne, J., & Dorr, C. (2014). Semantic plasticity and speech reports. *The Philosophical Review*, 123(3), 281–338.
- Hoek, D. (2018). Conversational exculpation. *The Philosophical Review*, 127(2), 151–196.
- Hyde, D. (2007). *Vagueness, logic and ontology*. Ashgate.
- Kennedy, C., & Stanley, J. (2009). On ‘Average’. *Mind*, 118(471), 583–646.
- King, J. (2014). The metasemantics of context-sensitivity. In A. Burgess & B. Sherman (Eds.), *New essays on Metasemantics*. Oxford University Press.
- Lasersohn, P. (1999). Pragmatic Halos. *Language*, 75, 522–551.
- Lauer, S. (2012). On the pragmatics of pragmatic slack. In A. Aguilar, R. Nouwen, & A. Chernilovskaya (Eds.), *Proceedings of Sinn and Bedeutung*, 16.
- Lewis, D. (1969). *Convention: A philosophical study*. Harvard University Press.
- Lewis, D. (1974). Radical Interpretation. *Synthese*, 23, 331–344.
- Lewis, D. (1975). Languages and language. In K. Gunderson (Ed.), *Minnesota studies in the philosophy of science* (Vol. VII, pp. 3–35). University of Minnesota Press.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(1), 339–359.
- Markman, E. (1989). *Categorization and naming in children: Problems of induction*. MIT Press.
- Murray, S. (2014). Varieties of update. *Semantics and Pragmatics*, 7, 1–53.
- Ninan, D. (2010). Semantics and the objects of assertion. *Linguistics and Philosophy*, 33(5), 355–380.
- Rabern, B. (2012). Against the identification of assertoric content with compositional value. *Synthese*, 189, 75–96.
- Raffman, D. (2014). *Unruly words*. Oxford University Press.
- Ripley, D. (2013). Sorting out the sorites. In F. Berto, E. Mares, F. Paoli, & K. Tanaka (Eds.), *Paraconsistency: Logic and applications*. Springer.
- Roberts, C. (2018). Speech acts in a discourse context. In D. Harris, & M. Moss (Eds.), *Daniel Fogal* (pp. 317–359). Oxford University Press.
- Shapiro, S. (2006). *Vagueness in context*. Oxford University Press.
- Soames, S. (1999). *Understanding truth*. Oxford University Press.
- Soames, S. (2002). *Beyond rigidity: The unfinished semantic agenda of naming and necessity*. Oxford University Press.
- Soames, S. (2009). The gap between meaning and assertion: Why what we literally say often differs from what our words literally mean. *Philosophical Essays*, 1, 278–297.
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Blackwell (2nd revised edition, 1995).
- Sperber, D., Clement, F., Heintz, C., Mascaro, O., Mercier, H., Origg, G., & Wilson, D. (2010). Epistemic vigilance. *Mind & Language*, 25(4), 359–393.
- Stalnaker, R. (1978). Assertion. *Syntax and Semantics*, 9, 315–332.
- Stalnaker, R. (1984). *Inquiry*. MIT Press.
- Stalnaker, R. (2014). *Context*. Oxford University Press.
- Stanley, J. (2005). Review of François Recanati, literal meaning. *Notre Dame Philosophical Reviews*.
- Stanley, J., & King, J. (2005). Semantics, pragmatics, and the role of semantic content. In Z. Szabó (Ed.), *Semantics versus pragmatics* (pp. 111–164). Oxford University Press.
- Stojnic, U. (2017a). On the connection between semantic content and the objects of assertion. *Philosophical Topics*, 45(2), 163–179.
- Stojnic, U. (2017b). Content in a dynamic context. *Nous*, 53(2), 394–432.
- Tappenden, J. (1993). The liar and sorites paradoxes: Toward a unified treatment. *Journal of Philosophy*, 90, 551–557.
- von Fintel, K., & Gillies, A. (2010). Must ... Stay ... Strong! *Nat Lang Semantics*, 18, 351–383.
- von Fintel, K., & Gillies, A. (2011). ‘Might’ made right. In A. Egan & B. Weatherson (Eds.), *Epistemic modality*. Oxford University Press.
- Wilson, D., & Sperber, D. (2002). Truthfulness and relevance. *Mind*, 111(443), 583–632.

- Yalcin, S. (2014). Semantics and metasemantics in the context of generative grammar. In A. Burgess & B. Sherman (Eds.), *Metasemantics* (pp. 17–54). Oxford University Press.
- Yli-Vakkuri, J. (2013). Propositions and compositionality. *Philosophical Perspectives*, 27, 526–563.

How to cite this article: Armstrong, J. (2023). Truth and imprecision. *Analytic Philosophy*, 00, 1–24. <https://doi.org/10.1111/phib.12286>