Context Dynamics *

Michael Caie
University of Toronto

Abstract In this paper, I consider how, given mutual knowledge of the information codified in a compositional semantic theory, an assertion of a sentence serves to update the shared information in a conversation. There is a standard account, due to Stalnaker, of how such conversational updating occurs. While this account has much to recommend it, in this paper I argue that it needs to be revised in light of certain patterns of updating that result from certain natural discourses. Having argued for this, I present a new account of conversational updating that can be seen as a natural generalization of the standard account, and show how it can predict these patterns in a simple and principled manner.

Keywords: context updating, Robert Stalnaker, pragmatics

There are two things called *contexts* that play important but distinct roles in standard accounts of language and communication. The first — call these *compositional contexts* — feature in a semantic theory. Compositional contexts are sequences of parameters that play a role in characterizing compositional semantic values for a given language, and in characterizing how such compositional semantic values determine a proposition expressed by a given sentence.[1] The second — call these *context sets* — feature in a pragmatic theory. Context sets are abstract representations of conversational states that

---

1 For a canonical discussion of such contexts see Kaplan 1989.

serve to determine the compositional contexts relevant for interpreting a speech-act and that such speech-acts act upon.[2]

In this paper, I'll consider how, given mutual knowledge of the information codified in a compositional semantic theory, an assertion of a sentence serves to update the context set. There is an orthodox account of how such conversational updating occurs.[3] However, while this account has much to recommend it, I'll argue that it needs to be revised in light of certain natural discourses.

In Section 1, I begin by outlining and motivating an account of conversational updating, that I call Standard Updating, that is similar to the account provided in Stalnaker 1978. Roughly speaking, according to this account, given a particular assertion, at each world in the context set there is a unique proposition expressed by this assertion. Updating, then, proceeds via *diagonalization*. To update the context set, we eliminate a world from this set just in case the unique proposition expressed at this world by the relevant assertion is false at this world.

In Section 2, I present an argument against Standard Updating. There I consider two natural discourses and argue that the proponent of Standard Updating cannot predict the patterns exhibited by these discourses in a principled and plausible manner.

In Section 3, I present an alternative account of conversational updating that I call Disjunctive Updating. According to the resulting picture, in asserting a sentence, at each world in the context set, we often express multiple propositions, which we update on, roughly, by eliminating a world just in case the disjunction of propositions expressed at that world is false at that world. According to this account, which propositions we express by the assertion of a sentence in a given discourse depends in a systematic way on which propositions have been expressed by other assertions within that discourse. I show how this account can predict the desired patterns of updating in a simple and principled manner.

Finally, in Section 4, I briefly compare Disjunctive Updating with three formally similar accounts of conversational updating proposed in Barker 2002, 2013, MacFarlane 2020 and King 2021 and offer some concluding remarks.

---

2 For a canonical discussion of such contexts see Stalnaker 1978.

3 See, for example, Stalnaker 1978, Stalnaker 2004, Stalnaker 1996, and Stalnaker 2014.

1 Standard updating

In Section 1.1, I introduce and motivate an orthodox account of conversational updating that I call Standard Updating, and, in Section 1.2, I defend the claim that this account can be seen as capturing the essential features of the account of conversational updating developed in Stalnaker 1978.

1.1 The account and its motivations

A compositional semantic theory for a language assigns semantic values to its lexical items, phrases and sentences, where the semantic value of a complex expression is determined as a function of the semantic values of its simpler constituents. Following Kaplan 1989, I'll assume that such semantic values are functions from pairs of compositional contexts and indices to extensions.

The elements of the compositional context serve to determine the (non-shiftable) values for context-sensitive items. These include indexicals and demonstratives, as well as other context-sensitive terms such as gradable adjectives. The elements of the index are those parameters that may be shifted by various operators.[4]

I'll assume that a sentence that is asserted has a syntactic structure in virtue of which compositional semantic values may be determined for the sentence and its syntactic constituents. In addition, I'll assume that, given such an interpreted sentence, a compositional context serves to determine a possible-worlds proposition. As a helpful abbreviation, we'll denote the intension associated with an expression ϕ, given a compositional context c: $[\![\phi]\!]^c$.[5]

---

4 According to Montague 1968, a compositional semantic value is a function from sequences of parameters to extensions. According to Kaplan 1989, following Kamp 1971 and Vlach 1973, we should distinguish between those parameters that may be shiftable by operators in the language — the elements of an *index* — and those that serve to interpret context-sensitive items of the language and are not so shiftable — the elements of the *context*.

5 It's worth noting that the assumption that a compositional context determines a proposition, given the semantic value of a sentence, is a slight idealization. To see this, consider the following simple picture of how this process of determination might work. Assume that the index contains a world parameter and for each non-world parameter p of the index, there is a well-defined notion of the value of p determined by the context c, which we'll write p_c. Then, given a compositional context, we can determine a possible-worlds proposition, given the semantic value of a sentence, by first saturating this function with a compositional context, and then saturating each of the non-world index parameters with the appropriate parameter

Following Stalnaker 1978, I'll assume that conversations take place against the background of a set of propositions that are mutually presupposed by the conversational participants. This is what we earlier called the *context set*, and what Stalnaker has also called the *common ground*. Presupposition may be thought of as a sui generis attitude that conversational participants adopt for the purposes of a conversation. I'll assume throughout that, in the cases of interest, the conversational participants adopt exactly the same presuppositions. We can represent the context set by a set of worlds, where a world is in the set just in case it's compatible with every proposition that's mutually presupposed by the conversational participants for the purpose of the conversation. I'll also assume, following Stalnaker 1978, that an assertion of a sentence can be thought of, roughly, as a suggestion to eliminate certain possibilities from the context set.

For a non-defective context set, I'll assume that it is mutually presupposed that the conversation is happening, what the presuppositions of the conversation are, and which assertions have been made in the conversation. Given an initial context set C, then, an assertion of a sentence ϕ immediately updates the context set by eliminating any worlds from C in which that as-

---

determined by the context, and, finally, lambda abstracting on the world parameter of the index. See Lewis 1980.

Unfortunately, this simple picture faces problems with assignment functions. While assignment functions are elements of the index—elements that are shifted by quantifiers—it isn't particularly plausible that the compositional context serves to pick out a unique such assignment function.

A better, if more complicated account, holds that a proposition is determined jointly by a compositional context — which fixes the values of the non-world, non-assignment function parameters of the index — and a set of assignments to the unbound occurrences of pronouns in a given sentence. Given such a set of assignments and a compositional context, we can determine a possible worlds proposition by employing essentially the same procedure as above, but, before lambda abstracting on the world parameter of the index, we first universally quantify over the variants of the assignment functions that agree with respect to the values that are determined for the unbound occurrences of pronouns.

To simplify our discussion, I'll suppress these complications and assume that a compositional context serves to determine a proposition for a sentence given its semantic value. To take account of assignment functions, however, we can employ the following translation procedure for the discussion that follows: talk about a world determining a compositional context or set of compositional contexts that interpret a given sentence may systematically be replaced with talk of a world determining a sequence or a set of sequences, consisting of a compositional context and an assignment of values to unbound occurrences of pronouns, that interpret a given sentence.

See Ninan 2010 for further discussion of different natural ways of determining something appropriate to playing the content role from compositional semantic values.

sertion of ϕ doesn't take place. We'll denote the context set that results from this update C_ϕ. Given an assertion of a sentence ϕ and a context set C, for each $w \in C_\phi$, there is, then, a corresponding token assertion of ϕ in w.

An assertion of a sentence, then, updates the context set on the event of the assertion of that sentence. Such an assertion, though, is intended to update the context set not just by adding the fact of the assertion as an additional presupposition, but by eliminating certain worlds from this set given facts about the context set and the semantic value of the sentence. Without prejudice to how exactly such possibilities are eliminated, we'll denote the result of this update to a context set C_ϕ: C^ϕ.

Given a particular assertion of ϕ in $w \in C_\phi$, the time, location and shared presuppositions of the conversation in which this assertion occurs are all fixed. I'll assume that these sorts of facts determine whether a compositional context *interprets* the particular assertion of ϕ in w. The interpretation of an assertion of a sentence by a compositional context is, as I'll understand it, a theoretical concept that should ultimately be understood in terms of the role that it plays in an account of conversational updating. Still, the concept is meant to have some intuitive content.

One gloss appeals to the concept of an assertion of a sentence *expressing* a proposition. According to this gloss, the compositional contexts that interpret an assertion of a sentence ϕ in w determine which possible-worlds propositions that assertion expresses in w.[6] In particular, when a compositional context c interprets an assertion of a sentence ϕ in w, then, in w, ϕ expresses $[\![\phi]\!]^c$.

According to a commonly accepted account of conversational updating, given a non-defective assertion of a sentence ϕ, for each world w in the context set, there will be a unique compositional context that interprets the corresponding occurrence of ϕ at w. According to this account, context sets are updated via diagonalization — that is, to update the context set we eliminate a world w from this class just in case w is incompatible with the proposition that ϕ expresses at w.

A little more precisely:

---

6 While I think that this is a useful gloss to bear in mind, and a useful way of talking that I'll employ throughout, it is also worth noting that the concept of a sentence expressing a proposition is one that likely has multiple precisifications of which this is just one. On this point, see the discussion in Section 1.2. For discussion of the distinction between the compositional semantic value role and the content of assertion role see Dummett 1959, Lewis 1980, Stanley 1997, Yalcin 2007, Ninan 2010, Rabern 2013, and Stonić 2017 amongst others.

Standard Updating: Given a non-defective assertion of a sentence ϕ in a conversation with a context set C, for each $w \in C_\phi$, there is a unique compositional context c_w that interprets the assertion of ϕ in w. The result of updating the context given the assertion is $C^\phi = \{w \in C_\phi : w \in [\![\phi]\!]^{c_w}\}$.

This account of conversational updating has much to recommend it.

First, it delivers the desired results in those cases in which there is a unique proposition expressed by an assertion of a sentence ϕ at each world in the context set, and each of these propositions agrees about the worlds in the context set. In such cases, given Standard Updating, the characteristic effect of an assertion of a sentence ϕ is to eliminate every world in the relevant context set incompatible with the unique proposition (restricted to the context set) expressed by ϕ at an arbitrary world in the context set.

Second, it delivers the desired results in those cases in which there is a unique proposition expressed by an assertion of a sentence ϕ at each world in the context set, but these propositions disagree about certain worlds in the context set.

To see this, consider a context set in which there is a world w_i in which the conversation is taking place at location l_i and a particular person, Ren, is at a distinct location l_j, and a world w_j in which the conversation is taking place at l_j and Ren is at l_j. Now suppose that the following sentence is asserted given a context set satisfying these constraints:

(R) Ren is here.

Given a standard semantics, if c_{w_i} and c_{w_j} are any compositional contexts that interpret (R) at w_i and w_j respectively, then we will have:

- $[\![(R)]\!]^{c_{w_i}} = \{w : \text{Ren is at } l_i \text{ in } w\}$

- $[\![(R)]\!]^{c_{w_j}} = \{w : \text{Ren is at } l_j \text{ in } w\}$

Since w_i is compatible with $[\![(R)]\!]^{c_{w_j}}$, but not compatible with $[\![(R)]\!]^{c_{w_i}}$, it follows that $[\![(R)]\!]^{c_{w_i}} \cap C_{(R)} \neq [\![(R)]\!]^{c_{w_j}} \cap C_{(R)}$. And so, in this case, there are distinct propositions expressed by the assertion of (R) at distinct worlds in the context set, and these propositions disagree over certain worlds in the context set.

Despite this fact, though, it is clear that one can non-defectively assert (R) when there is contextual uncertainty about the location of Ren. Indeed,

it is clear what the updating effect of an assertion of (R) on this context set would be; it would eliminate all the worlds in this set in which Ren is not at the location of utterance in that world. This is exactly the verdict delivered by Standard Updating.

1.2   Stalnaker and standard updating

Standard Updating corresponds roughly to the account of conversational updating that can be found in Stalnaker 1978.[7] There are, however, a few interpretive issues concerning this account that are worth briefly discussing.

One difference between Standard Updating and the account proposed in Stalnaker 1978 concerns the explicit assumption that there is a unique compositional context that interprets a given assertion for each world in the context set. While Standard Updating imposes this constraint, Stalnaker (1978) does not explicitly do so. Nonetheless, while this constraint is not explicitly imposed, it seems clear to me that Stalnaker, and subsequent authors who endorse roughly the same pragmatic framework, tacitly assume such uniqueness.

To see this, note that if there may be multiple compositional contexts that interpret a given assertion of a sentence ϕ at some world w in the context set, then, assuming that updating proceeds locally by diagonalization, the question arises under what conditions should w be eliminated from the context set given that the assertion of ϕ is interpreted by multiple compositional contexts at w? And the answer to this question is not obvious. For example, one might maintain that w is eliminated if *all* of the propositions determined by some compositional context that interprets ϕ at w are false at w, or one might maintain that w is eliminated if *some* of the propositions determined by some compositional context that interprets ϕ at w are false at w, or one might have some other account. The key point is that while it isn't obvious what the answer to this question should be, this question does not arise in Stalnaker 1978, nor is it the sort of question that has been addressed by most authors who accept Stalnaker's pragmatic account of conversational updating. The reason for this, I take it, is just that it is tacitly assumed that, given a non-defective assertion of a sentence ϕ, there is a unique compositional context that interprets an assertion of ϕ at each world in the relevant context set, and, given this, it is clear how updating by diagonalization should proceed.

---

7 For a recent critical discussion of this model of updating see Kirk-Giannini 2020.

Here's another way in which the present account may seem to differ from the account proposed in Stalnaker 1978. There are some remarks in Stalnaker 1978 that suggest that, given a non-defective assertion of a sentence relative to a context set C, for all worlds $w, w' \in C$, any proposition expressed by the sentence in w or w' must agree about their truth-values at worlds in C. Given this constraint, we can speak of there being a single proposition (restricted to C) that is expressed by the sentence ϕ at each world in the context set. On this picture, the characteristic effect of an assertion of a sentence ϕ is to eliminate every world in the initial context set incompatible with this unique proposition.

It may initially seem that this marks a significant difference between Standard Updating and the account proposed in Stalnaker 1978. However, I think that once certain distinctions are drawn it is clear that this feature of Stalnaker's account is, in fact, compatible with Standard Updating.

To see why it might seem like the account of updating presented in Stalnaker 1978 has roughly the form just sketched, first note that Stalnaker there claims that a basic constraint on a non-defective assertion is that "[t]he same proposition is expressed relative to each possible world in the context set." In a footnote he clarifies that *sameness* here means that the propositions are the same when restricted to the context set. He then notes that there would seem to be violations of this constraint. In response, he suggests that we take the proposition expressed by ϕ to be the *diagonal proposition*, which we can take to be: $\{w_i \in C_\phi : w_i \in [\![\phi]\!]^{c_{w_i}}\}$.

To clarify what exactly is going on here and how it relates to our preceding discussion, let's distinguish between two senses in which an assertion of a sentence may be said to *express* a proposition. First, let's say that a sentence ϕ expresses$_1$ a proposition p in w just in case $p = [\![\phi]\!]^c$, for some compositional context c that interprets ϕ in w. This is the notion of expression that we've been working with so far. It is, moreover, a notion of expression that Stalnaker should countenance, since it is in terms of this relation that the diagonal proposition is defined. Second, let's say that a sentence ϕ expresses$_2$ a proposition p in w just in case p is the result of successfully updating the context set given the assertion of ϕ in w.

Now it should be clear that it is simply not a reasonable constraint on a non-defective assertion that the same proposition is expressed$_1$ relative to each possible world in the context set. For, as we've seen, this constraint isn't satisfied in many cases of non-defective assertions, such as the assertion of

(R) considered above. It would be a mistake, then, to see Stalnaker as trying to vindicate this principle.

It is, however, a reasonable constraint on a non-defective assertion that the same proposition is expressed$_2$ relative to each possible world in the context set updated on the event of the assertion. For this tells us that, for a non-defective assertion, the suggested update is the same, for each world in this set. But if this were violated, it would be unclear what the suggested update of the assertion is. And this, indeed, would be a communicative defect.

This constraint, moreover, is satisfied given Standard Updating. I suggest, then, that Stalnaker (1978) should be seen as endorsing this latter constraint, and as proposing an account of updating at least quite close to Standard Updating as a way of satisfying it.

2 Against standard updating

Despite its many virtues, I think that this account of conversational updating should be rejected. In this section, I'll present an argument against Standard Updating.

In Section 2.1, I'll consider two natural discourses, and in Section 2.2, I'll argue that the proponent of Standard Updating cannot predict the patterns of updating exhibited by these discourses in a principled way. In Secion 3, I'll show that, given the assumption that there may be multiple compositional contexts that interpret a given assertion in a world, there is a natural account of conversational updating that predicts the patterns exhibited in these discourses.

The lesson to be drawn from this, I'll conclude, is that Standard Updating is flawed because it encodes a false uniqueness assumption. In order to account for the distinctive pattens of updating exhibited by certain discourses, in a principled manner, we must allow that there may be multiple compositional contexts that interpret an assertion of a sentence at a given world in a context set.

2.1 Sarah's socks

Consider the following two discourses.

Minimal Presuppositions: There is a conversation amongst a group of Sarah's friends in which the following facts are all presupposed:

- Sarah has a collection of individual socks that she wears in various combinations. She currently has four individual socks — two are striped and two are solid.

- Her friends all refer to pairs whose member socks are both striped or both solid as 'matching', and pairs whose members are respectively striped and solid as 'mixed'.

- Another one of their friends, Tim, has strong preferences about pairs of socks. Either: (i) he likes matching pairs and dislikes mixed ones or, (ii) vice versa, but their presuppositions are silent about which.

The Facts: In addition, the following facts obtain:

- Tim likes matching socks and dislikes mixed ones.

A speaker knows that the Facts obtain and wants to communicate this to the other conversational participants.

Given the listed presuppositions, I submit that each of the following discourses provides a reasonable way, available to the speaker, of communicating that the Facts obtain.

Tim Likes Matching:	**Tim Dislikes Mixed:**
(1) Sarah has two pairs of socks.	(1) Sarah has two pairs of socks.
(2) Tim likes both of them.	(4) Tim dislikes both of them.
(3) Both of them are matching.	(5) Both of them are mixed.

The first point to note about these discourses is that, on the intended reading of the first sentence, the items in Sarah's wardrobe that are in the domain of quantification and satisfy the predicate 'pair of socks' consist of *non-overlapping* combinations of the individual socks that Sarah owns. This is the reading that would allow the speaker in either discourse to continue by truly uttering, for example, 'You could wear the first pair, and I could wear the other.'[8]

---

8 The predicate 'pair of socks' belongs to a class of expressions we can call *configurational predicates*. Examples of such predicates also include: 'outfit', 'dining set', 'team', 'hand', in

The next point to note is that there are two classes of interpretations that satisfy this constraint. According to a *mixed interpretation*, the combinations of Sarah's socks that are in the domain of quantification and satisfy 'pair of socks' consist of pairs one member of which is striped and one member of which is solid. According to a *matching interpretation*, the combinations of Sarah's socks that are in the domain of quantification and satisfy 'pair of socks' consist of pairs such that both members are striped or both members are solid.

To see, intuitively, how these discourses serve to communicate the desired information about Tim's preferences, let us focus on Tim Likes Matching — the same sort of reasoning may be applied, mutatis mutandis, in Tim Disikes Mixed.

Initially, we can divide the worlds in the context set into two classes depending on Tim's preferences — there are those worlds in which Tim likes mixed pairs of socks and those worlds in which Tim likes matching pairs of socks.

Now the first sentence in this discourse doesn't have any informational effect. It does, however, serve to cue us in to a class of available interpretations. For each world in the context set, given the assertion of (1), there are two types of available interpretations — mixed interpretations and matching interpretations.

The assertion of the second sentence also provides no information about Tim's preferences. The assertion of this sentence does, however, serve to constrain which interpretations are available, given Tim's preferences. For, if Tim likes mixed pairs, then the assertion of 'Tim likes both of them' is true only under a mixed interpretation, while if Tim likes matching pairs then the assertion of this sentence is true only under a matching interpretation. Given the assertion of this sentence, then, at worlds at which Tim likes mixed pairs the only remaining available interpretations are mixed, while at worlds in

---

the context of card games, or certain occurrences of 'word', for example in the context of Scrabble. Krifka 2009 provides a detailed semantic proposal for such expressions. These sorts of predicates are briefly mentioned in Kratzer 2012. What unites this class is that, on their natural interpretations, such predicates may be satisfied by various combinations of more basic objects — for example, combinations of individual socks, combinations of various items of clothing, combinations of forks, knives, spoons, or combinations of individuals. In the present context, configurational predicates are interesting because they provide simple and tractable cases in which there are clear symmetries that plausibly impose a lower bound, greater than one, on the number of compositional contexts that interpret a given occurrence of a certain sentence in a given world.

which Tim likes matching pairs the only remaining available interpretations are matching.

Given this, then, it follows that the third assertion will be false at each world in the context set at which Tim likes mixed pairs, given the available interpretations of this sentence at that world, while the third assertion will be true at each world in the context set at which Tim likes matching pairs, given the available interpretations of this sentence at that world.

It follows, then, that, given the preceding two sentences, the third assertion serves to rule out all of the worlds in which Tim likes mixed pairs, leaving only worlds in the context set in which Tim likes matching pairs.

That, at any rate, is one way of seeing how the relevant information is conveyed. Predicting and vindicating this reasoning, given a general account of conversational updating, will have to wait until Section 3.

Now I take it that, in addition to the presuppositions listed in the description of Sarah's Socks, in assessing the updating effects of these discourses, it is tacitly assumed that the conversational participants make certain reasonable semantic presuppositions. Indeed, we appealed to such additional presuppositions in the preceding explanation. For the argument that follows, it will be helpful to make these additional presuppositions explicit.

The sentence 'Sarah has two pairs of socks.' is clearly context sensitive. For while this has a true reading, given the facts about Sarah's wardrobe, there are also contexts in which, given the same facts, one could truly utter 'Sarah has six pairs of socks.' The latter, for example, has a natural true reading if we're interested in the number of different ways in which we could pair her socks.

There are two natural accounts of the locus of this context sensitivity. According to the first, 'pair of socks' itself is context-sensitive. According to the second, the relevant context sensitivity consists in a contextually variable restriction on the numerical quantifiers that occur in this and related sentences.[9] I'm inclined to think that the latter is the more plausible of these two options. In what follows, however, I'll remain neutral about the relevant source of context sensitivity for the sentences in our discourses.

A compositional context c determines a function that maps a world w to the class of items that are in the domain of quantification at w given c, are owned by Sarah in w, and satisfy 'pair of socks' at w given c.

---

[9] See, for example, von Fintel 1994. In this case, I think the most plausible account will postulate a covert variable that is interpreted by context and that restricts the quantifiers by intersecting with the interpretation of 'pair of socks'. See Stanley & Gendler Szabó 2000.

Def: We'll say that such a function is a *dressing intension* just in case:

(i) The value of this function, for each world $w \in C_{(1)}$, contains exactly two pairings of socks.

(ii) These pairings are non-overlapping.

(iii) The value of this function is the same, for any two worlds $w, w' \in C_{(1)}$.

As noted above, we can partition the class of dressing intensions into two sub-classes.

Def: We'll say that a dressing intension is *mixed* just in case its extension across the worlds in $C_{(1)}$ contains only mixed pairs of Sarah's socks.

Def: We'll say that a dressing intension is *matching* just in case its extension across the worlds in $C_{(1)}$ contains only matching pairs of Sarah's socks.

In addition to the presuppositions listed in the description of Sarah's Socks, let us *stipulate* that the conversational participants presuppose that the following all hold:

- There is at least one compositional context that interprets each assertion in any relevant occurrence of Tim Likes Matching and Tim Dislikes Mixed.

- Any compositional context that interprets a sentence in one of these discourses determines a dressing intension.

- If a compositional context c interprets the occurrence of (2)–(5), in either discourse, then, for each world w, the class of items that are in the domain of quantification at w given c and satisfy 'them' at w given c is exactly the class of items that are in the domain of quantification at w given c, are owned by Sarah in w, and satisfy 'pair of socks' at w given c.

- Subject to the preceding constraints, if a compositional context c interprets an occurrence of (1)–(5) then it assigns the obvious semantic values to lexical items in this sentence, and satisfies the obvious

constraints on how the semantic values are determined for complex expressions.

Henceforth, we'll take *the Presuppositions* to consist of these propositions as well as those initially listed, and we'll let C be the set of worlds in which the Presuppositions hold. While such presuppositions may not be required in order for the above discourses to give rise to the relevant patterns of updating, I take it that such presuppositions may in principle be adopted, and that, given such presuppositions, these discourses do result in the relevant patterns of updating.

2.2 An argument against standard updating

In this section, I'll present an argument for rejecting Standard Updating. In the preceding section, we considered two discourses and noted that each discourse provides a natural way of communicating the facts about Tim's preferences. In this section, I'll argue that the proponent of Standard Updating cannot predict the patterns of updating exhibited by these discourses in a principled manner.

There are two notable features of these discourses that any adequate account of updating should be able to predict. The first feature is simply that, given the Presuppositions, each of these discourses provides a way of communicating the Facts. The second feature is that, given the Presuppositions and the Facts, these discourses are not only able to communicate this information, but are guaranteed to do so without thereby communicating any false information.

These both strike me as being manifest features of these discourses. To maintain otherwise would be to either deny that these discourses impart the information that they do, or to impute a certain riskiness to these discourses that just doesn't seem to be present.

We can codify both of these features in the following principle:

> **Safe Information:** Letting $F \subset C$ be the set of worlds in which the Presuppositions and the Facts hold, the following two conditions are satisfied:
>
> (i) $C^{(1)(2)(3)} \subseteq F$ and $C^{(1)(4)(5)} \subseteq F$.
>
> (ii) For each $w \in F$, if Tim Likes Matching occurs at w, then $w \in C^{(1)(2)(3)}$, and if Tim Dislikes Mixed occurs at w, then $w \in C^{(1)(4)(5)}$.

This principle provides a precise sense in which these discourses are guaranteed to communicate the relevant information without thereby communicating any falsehood. Condition (i) tells us that these discourses are guaranteed to communicate the relevant information, given the Presuppositions, while condition (ii) tells us that each discourse is guaranteed to not communicate any false information, given the Presuppositions and the Facts.

I'll now argue that the proponent of Standard Updating can't vindicate Safe Information in a plausible and principled manner. The argument for this will take the following form. First, I'll argue for a certain principle — Minimal Symmetry — that codifies a minimal sense in which (at least some of) the same interpretations are available for the common first sentence in both discourses. Assuming this principle, I'll argue that, given Standard Updating, at least one of the conditions imposed by Safe Information will fail.

The argument for this conclusion will take the form of a dilemma. To this end, I'll consider a second plausible principle — Preservation — that tells us that the interpretations of the second sentences in our discourses are inherited from the interpretation of the first sentence.

I'll first show that, given Standard Updating, Safe Information must fail given Minimal Symmetry and Preservation. In particular, given Standard Updating and these two principles, it follows that it is possible, compatible with the Facts and the Presuppositions, for each of our discourses to result in the adoption of some false presuppositions.

I'll then consider the prospects for avoiding this failure of Safe Information by rejecting Preservation. Roughly speaking, the proponent of Standard Updating can endorse Minimal Symmetry without thereby predicting that our discourses may result in the adoption of some false presuppositions by appealing to a certain type of context shift. I'll argue, however, that the only plausible general principle that predicts the desired context shifts also leads to the prediction that these discourses must fail to result in the Facts being presupposed.

Now I'm inclined to think that Preservation is quite plausible, and in the following section I'll show how, by rejecting Standard Updating, one can vindicate this principle as well as Minimal Symmetry and Safe Information. The claim in this section, however, is that, regardless of whether the proponent of Standard Updating endorses Preservation, there is no principled way for them to vindicate Safe Information.

To begin, then, consider the following principle:

Minimal Symmetry: Again letting $F \subset C$ be the set of worlds in which the Presuppositions and the Facts hold, there are at least two worlds $w_1, w_2 \in F$ such that:

(i) Tim Likes Matching occurs at w_1 and Tim Dislikes Mixed occurs at w_2.

(ii) There is a compositional context c_{w_1} that interprets (1) at w_1, and a compositional context c_{w_2} that interprets (1) at w_2, such that c_{w_1} and c_{w_2} determine the same dressing intension.

This principle tells us that there are at least two worlds compatible with the Presuppositions and the Facts that differ with respect to which of our two discourses occur, but that agree with respect to the interpretation of their common first sentence.

To see why this is plausible, note that it is plausible that the interpretation of the first sentence of these discourses should be determined by the facts that obtain up to and including the time of the assertion of this sentence. But it is quite plausible that there are worlds compatible with the Presuppositions and the Facts that are exactly alike up to and including the time of the assertion of (1) and that differ only afterwards in how the relevant discourse proceeds. We can codify these claims in the following principles that jointly entail Minimal Symmetry:

Supervenience: Amongst the worlds in which the Presuppositions hold, any two worlds, w_1 and w_2, that agree about all of the facts that obtain up to and including the time of the assertion of (1) in either Tim Likes Matching or Tim Dislikes Mixed will be such that for each compositional context c_1 that interprets the assertion of (1) in w_1 there is a corresponding compositional context c_2 that interprets the assertion of (1) in w_2 that agrees with c_1 about the interpretation of (1) and its constituents.

Metasemantic Symmetry: Amongst the worlds in which the Presuppositions and the Facts hold there is some world w_1 in which Tim Likes Matching occurs and some world w_2 in which Tim Dislikes Mixed occurs that agree about all of the facts that obtain up to and including the time of the assertion of (1).

One could reject Supervenience by maintaining that how (1) is interpreted depends, in part, on which sentences are asserted *after* the assertion of (1). In particular, one might maintain that amongst the worlds in which the Facts obtain, at worlds at which the assertion of (1) is followed by (2), the occurrence of (1) is interpreted by a compositional context that determines a matching dressing intension, while, at worlds at which (1) is followed by (4), the occurrence of (1) is interpreted by a compositional context that determines a mixed dressing intension.

I'm inclined to think, however, that it is quite implausible that how an assertion of this sentence at some time t is interpreted depends on what sentences are asserted after t. For example, Tim Likes Matching and Tim Dislikes Mixed may each take place over an extended period time, and there may be long gaps — in principle, months or years — between each assertion. You might imagine either of these discourses occurring over text message with the predictable lags between each assertion. It seems to me implausible to maintain that, in such a case, it is only in virtue of the fact that, say, (2) is asserted a month after the initial assertion of (1), that (1) expresses the proposition(s) that it does.[10]

To see why Metasemantic Symmetry is plausible, consider the following situation:

> A speaker, Sam, wants to communicate the Facts, in a communicative situation in which the Presuppositions obtain. Sam knows that he can assert either the sentences in Tim Likes Matching or the sentences in Tim Dislikes Mixed to achieve his communicative goals. He begins by asserting (1). Sam, however, has no views about which combinations of socks count as being in the domain of quantification and satisfying 'pair of socks', and has no intentions about whether to continue by asserting (2) and (3) or by asserting (4) and (5). Indeed, the underlying facts about Sam make it a genuinely chancy matter how he will continue, it being equally likely that he will continue either way.

---

10 It's worth noting that this sort of view has been defended. See, for example, Jackman 1999. The proponent of this view may, then, take the following argument to provide support for their view. However, we'll see in the next section that if we give up the assumption that there is a unique compositional context that interprets a given assertion at a world w, then we can predict the patterns of updating in our discourses, while endorsing both Supervenience and Metasemantic Symmetry. Since I think these principles are more plausible than this uniqueness assumption, the lesson that I'm inclined to draw from the following argument is not that we should reject Supervenience but instead that we should reject this uniqueness assumption.

It would certainly seem that there will be worlds in $C_{(1)}$ compatible with the Facts in which the speaker satisfies these conditions. Moreover, given the chance facts, such worlds will come in pairs. These worlds will be perfect duplicates up and including the time of the assertion of (1) and will differ thereafter. In one of these worlds the chancy process of continuation results in Tim Likes Matching occurring, while in the other this process results in Tim Dislikes Mixed occurring. Since these two worlds are duplicates up to and including the time of the assertion of (1) though, they will agree about all of the facts that obtain up to and including the time of the assertion of (1) and so will witness the truth of Metasemantic Symmetry.

Minimal Symmetry, then, strikes me as being quite plausible. In what follows, I'll take this principle for granted. Given this background assumption, I'll now argue that the proponent of Standard Updating can't predict the truth of Safe Information in a principled manner.

Consider next the following principle:

> **Preservation:** Given that the Presuppositions obtain, for each $w \in C_{(1)}$ if there is a unique compositional context c_w that interprets an assertion of (1) in w, then if $w \in C^{(1)}$, then c_w interprets the subsequent assertion of either (2) or (4) in w.

This principle may be motivated by appeal to the claim that, on their intended readings, in each discourse, all of the sentences are jointly interpreted in a uniform manner, so that the domains of quantification remain the same, and the interpretations of the pronouns are inherited in the natural manner from the interpretations of the preceding sentences. This strikes me as being very plausible. For when we consider, in each discourse, which things the speaker is claiming Tim likes, the natural answer is the things that they were talking about in the preceding sentence.

While I take it that both Minimal Symmetry and Preservation are quite plausible, we can show that, given Standard Updating, these principles entail the falsity of Safe Information.

> **Claim:** Given Standard Updating, Safe Information is inconsistent with Minimal Symmetry and Preservation.

> Safe Information requires that, for each $w \in F$, if Tim Likes Matching occurs at w, then $w \in C^{(1)(2)}$, and if Tim Dislikes

Mixed occurs at w, then $w \in C^{(1)(4)}$. We can show, however, that given Standard Updating, Minimal Symmetry and Preservation at least one of these conditions must fail to hold.

To see why, note that Standard Updating and Minimal Symmetry entail that there are at least two worlds $w_1, w_2 \in F$ such that Tim Likes Matching occurs at w_1, Tim Dislikes Mixed occurs at w_2, and there are unique compositional contexts c_{w_1} and c_{w_2} that determine the same dressing intension that each respectively uniquely interpret the occurrence of (1) in each of these discourses.

There are, then, two possibilities:

(a) c_{w_1} and c_{w_2} both determine a mixed dressing intension.

(b) c_{w_1} and c_{w_2} both determine a matching dressing intension.

Given Preservation, though, both (a) and (b) are incompatible with Safe Information.

If (a) is the case, then, given Preservation, it follows that (2) will be interpreted so that it is true at those worlds in which Tim likes mixed pairs of socks. However, since w_2 is a world in which Tim likes matching pairs of socks and dislikes mixed pairs, it follows that $w_2 \notin C^{(1)(2)}$.

If (b) is the case, then, given Preservation, it follows that (4) will be interpreted so that it is true at those worlds in which Tim dislikes matching pairs of socks. But since w_1 is also a world in which Tim likes matching pairs of socks and dislikes mixed pairs, it follows that $w_1 \notin C^{(1)(4)}$.

The above result shows that, assuming the truth of Minimal Symmetry and Preservation, the proponent of Standard Updating incorrectly predicts the failure of Safe Information. Since I think that Minimal Symmetry and Preservation are both true, I think that this result suffices to show that the proponent of Standard Updating cannot adequately predict the patterns of updating witnessed by our discourses.

It is, however, worthwhile considering whether the proponent of Standard Updating might be able to endorse Safe Information, in a principled way, by rejecting Preservation. I'll argue now that they cannot.

The preceding failure of Safe Information stemmed from the fact that, given Minimal Symmetry, it follows that, for at least one of our discourses, there is some world in the context set and some interpretation of (1) at that world that, if passed along to the succeeding sentence in this discourse, results in that sentence expressing a false proposition at this world. And Preservation ensures that the relevant interpretation is passed along to the succeeding sentence, and so we get a failure of Safe Information.

A natural way to resist this conclusion is to maintain that, in the problematic cases, the undesirable interpretation doesn't get passed along. And indeed, for both discourses, there is some interpretation that is compatible with the Presuppositions that ensures that the second sentence in this discourse expresses a truth, given the Facts. For example, suppose that Tim Likes Matching occurs at $w_1 \in$ F and c_{w_1} — which uniquely interprets (1) at w_1 — determines a mixed dressing intension. Given this, it follows from Preservation, that (2) is false at w_1. However, there is another interpretation that is compatible with the Presuppositions that would make (2) true, namely one that determines a matching dressing intension. The same point applies, mutatis mutandis, for Tim Dislikes Mixed.

While we can block the above argument by appealing to this sort of context shift, I take it that if this is to provide a principled response on behalf of the proponent of Standard Updating there should be some plausible general principle from which it follows.

Consider, then, the following principle:

> **Uniform Charity:** For each world $w \in$ C, if ϕ is a sentence asserted in one of Tim Likes Matching or Tim Dislikes Mixed in w, then if there is some compositional context c that satisfies the conditions imposed by the Presuppositions and is such that $w \in [\![\phi]\!]^c$, then any compositional context c' that interprets the assertion of ϕ at w satisfies the conditions imposed by the Presuppositions and is such that $w \in [\![\phi]\!]^{c'}$.

The principle codifies a general preference for those interpretations, compatible with the Presuppositions, that make an assertion in our discourses true over those that make the assertion false.

This principle entails the desired context shift. Moreover, so far as I can see, appealing to Uniform Charity is the only reasonably principled way of predicting the desired shift in context. For, so far as I can see, the only relevant fact about the putative shift in context is that it is exactly the shift that's

required in order for the assertion of the second sentence in the relevant discourse to express a true proposition at the world in which it is asserted.

Appealing to this principle, however, does not allow the proponent of Standard Updating to endorse Safe Information. For we can show that Standard Updating and Uniform Charity are jointly inconsistent with Safe Information.

> **Claim:** Standard Updating and Uniform Charity entail the falsity of Safe Information.

Safe Information requires that, given the Presuppositions, updating on either Tim Likes Matching or Tim Dislikes Mixed rules out all the worlds in the context set that in which the Facts don't obtain.

Given Standard Updating and Uniform Charity, though, the result of updating C on Tim Likes Matching will simply be the set of worlds in C in which this discourse occurs, and similarly for the result of updating C on Tim Dislikes Mixed. And both sets will contain worlds in which Tim likes mixed pairs of socks and worlds in which Tim likes matching pairs of socks, since such worlds are compatible with the Presuppositions.

To see why this is so, first note that, for each $w \in C$, and each of (1)–(5), there is some interpretation compatible with the Presuppositions that ensures that the proposition expressed by that sentence is true at w. On the one hand, if, at w, Tim likes matching pairs and dislikes mixed ones, then (1), (2) and (3) are all true under a compositional context that determines a matching dressing intension, while (1), (4) and (5) are all true under a compositional context that determines a mixed dressing intension. On the other hand, if, at w, Tim likes mixed pairs and dislikes matching ones, then (1), (2) and (5) are all true under a compositional context that determines a mixed dressing intension, while (1), (3) and (4) are all true under a compositional context that determines a matching dressing intension.

Now, updating C on Tim Likes Matching automatically eliminates every world in C in which this discourse doesn't take place. However since, for every $w \in C$ and every sentence in Tim Likes Matching, there is some interpretation of that

sentence, compatible with the Presuppositions, such that the resulting proposition is true at w, it follows, given Standard Updating and Uniform Charity, that updating C on Tim Likes Matching will *only* eliminate those worlds in C in which this discourse doesn't take place.

To see this, let w be some world in C in which Tim Likes Matching occurs. Since, for each sentence in this discourse there is some interpretation, compatible with the Presuppositions, that makes the relevant assertion true at w, Standard Updating and Uniform Charity require that the relevant sentence in w be interpreted by such a compositional context. It follows that w will remain uneliminated, given the assertion of each of the sentences in Tim Likes Matching. Thus, given Standard Updating and Uniform Charity, the result of updating C on Tim Likes Matching will simply be the set of worlds in C in which this discourse occurs. And since the resulting set contains worlds in which Tim likes mixed pairs of socks and worlds in which Tim likes matching pairs of socks, we have a failure of condition (i) of Safe Information.

The same points apply, mutatis mutandis, to the result of updating C on Tim Dislikes Mixed.

I've argued that the proponent of Standard Updating cannot vindicate the truth of Safe Information in a principled manner. On the one hand, if they endorse Preservation, then they incorrectly predict that there are failures of condition (ii) of Safe Information. On the other, if they are to avoid predicting this type of failure of Safe Information, then they must postulate a certain context shift. I've argued though that appealing to such context shifts in a principled and uniform manner means that the proponent of Standard Updating incorrectly predicts that there are failures of condition (i) of Safe Information. I conclude that the proponent of Standard Updating can't accommodate the patterns of updating exhibited by our two discourses in a plausible and principled manner.

In the next section, we'll see how one can predict in a simple way the truth of Safe Information by giving up the assumption that, for a given assertion in some world, there is at most one compositional context that interprets that assertion at that world. Notably, this account is jointly compatible with

Preservation, Supervenience and Metasemantic Symmetry (and so, also, Minimal Symmetry).

3 Disjunctive updating

A number of authors have noted that there are many cases of context-sensitive terms where it is at least prima facie implausible that there are, in general, features of the situation in which a given utterance takes place that determine a unique resolution of the relevant compositional contextual parameters.[11]

In this section, we'll bolster the case for there being multiple compositional contexts that interpret a given assertion in a world by showing that, given this hypothesis, there is a natural account of conversational updating that predicts the patterns exhibited by our discourses.

According to this account, given a non-defective assertion of a sentence ϕ, there will be, for each world w in the relevant context set, a non-empty set of compositional contexts that interpret that assertion of ϕ at w. Given such an assertion, the context set is updated by eliminating a world w from this class just in case w is incompatible with *every* proposition that ϕ expresses at w.

A little more precisely:

> **Disjunctive Multi-Context Updating:** Given a non-defective assertion of a sentence ϕ in a conversation with a context set C, for each $w \in C_\phi$, there is a non-empty set of compositional contexts that interpret ϕ in w: I_w^ϕ. The result of updating the context given the assertion is $C^\phi = \{w \in C_\phi : w \in [\![\phi]\!]^c$, for some $c \in I_w^\phi\}$.

According to this account, a world w is eliminated from a context set, given an assertion, only when the *disjunction* of those propositions that are

---

11 For arguments that our communicative intentions together with other relevant facts often leave the resolution of certain context sensitive terms underdetermined see, for example, Braun & Sider 2007, Egan 2009, Buchanan 2010, von Fintel & Gillies 2011, King 2014, King 2018, MacFarlane 2020, and King 2021. Many arguments for this claim have focussed on gradable adjectives such as 'tall', 'heavy', 'happy', 'fast'. For standard accounts of the semantics of such expressions see, for example, von Stechow 1984, Kennedy 1999, and Kennedy 2007. See Barker 2002 and MacFarlane 2020 for accounts of updating with expressions involving gradable adjectives that are similar the account that follows.

expressed by that assertion at w is false at w. In those cases in which there is a unique compositional context that interprets a given assertion at each world in the context set, this account delivers the same results as Standard Updating. In a certain sense, then, we can think of this account as a natural generalization of the former account shorn of the requirement that there be a unique compositional context that interprets a non-defective assertion at a given world in the context set.

In addition to the general account of updating codified in Disjunctive Multi-Context Updating, I'll assume that any world in a non-defective context set satisfies a certain constraint on the compositional contexts that interpret the assertions within a given discourse. In particular, I will assume that each world in a non-defective context set is such that if the assertion of a sentence β immediately follows the assertion of a sentence α within a discourse in that world, then if α expresses at that world some propositions that are true at that world, then the compositional contexts that interpret the subsequent assertion of β at that world are exactly those that determine the true propositions expressed by α.

A little more precisely, I'll assume:

> **Contextual Pruning:** If the assertion of a sentence β immediately follows the assertion of a sentence α within a discourse in some world w, then if $\{c : c \in I_w^\alpha$ and $w \in [\![\alpha]\!]^c\} \neq \varnothing$, then $I_w^\beta = \{c : c \in I_w^\alpha$ and $w \in [\![\alpha]\!]^c\}$.

To have a useful label, we'll call the conjunction of Disjunctive Multi-Context Updating and Contextual Pruning: Disjunctive Updating.

This account of conversational updating provides the following picture of how a context set evolves throughout a given discourse. First, given the initial assertion in the discourse, for each world in the relevant context set, there is a non-empty set of compositional contexts that interpret this assertion at that world. The context set is then updated by removing all and only the worlds where this assertion does not express some true proposition. Given a subsequent assertion, it is then determined which compositional contexts interpret that assertion at each world in the updated context set in which this assertion occurs — it's the subset of those compositional contexts that determined a true proposition for the previous assertion. This, then, determines how the context set is updated, given this assertion, as well as which compositional contexts interpret the next sentence that is asserted in the discourse. And so on.

On the resulting picture, then, we can think of an assertion as acting on the context set in three distinct ways. First, an assertion cuts down on the worlds in the context set by removing all those worlds in which such an assertion does not take place. Second, an assertion cuts down on the remaining worlds in the context set by removing all those worlds in which that assertion expresses only false propositions. Third, an assertion cuts down on the number of compositional contexts that are available to interpret a subsequent assertion.

It's worth noting that there's a natural way of modelling the process of conversational updating that delivers exactly the results predicted by Disjunctive Updating.

Call a *context fragment* an ordered pair of a compositional context and a world, and call the *fragmentation* of C_ϕ the set of context fragments $\langle c, w \rangle$ such that $w \in C_\phi$ and c interprets ϕ in w. Let C be the initial context set in which a discourse occurs and let α be the initial sentence asserted in this discourse. Then we'll call the fragmentation of C_α the set of *initial context fragments* for that discourse.

According to this picture, to determine the updating effects of a given discourse we begin with the set of initial context fragments for that discourse. An assertion of a sentence ϕ in the discourse first eliminates those fragments in which that assertion doesn't take place at the world parameter of the discourse, and second eliminates those remaining fragments in which the proposition expressed by ϕ, given the compositional context parameter, is false at the world parameter. At each stage of this process, we can recover the updated context set as the set of worlds that occupy the second position in some member of the resultant set of context fragments. It can be verified that this process will generate the same patterns of updating as Disjunctive Updating.[12]

Let's now consider how the proponent of this account can predict the patterns of updating exhibited in Tim Likes Matching and Tim Dislikes Mixed. We again let C be the set of worlds in which the Presuppositions hold. Call

---

12 It is suggested in Barker 2002, 2013 that a conversational state may be modelled by appeal to a set of ordered-pairs roughly like this, and that such conversational states may be updated in roughly this manner. MacFarlane 2020 also maintains that we can model conversational updating by appeal to the sorts of sets of ordered-pairs that Barker proposes. I discuss the relation between these accounts of updating and the one presented here in Section 4. For some other accounts with some similarities to these accounts see van den Berg 1996, Lederman 2014 Chapter 2, Muñoz 2020 and Kocurek, Jerzak & Rudolph 2020.

The Disjunctive Account the conjunction of Disjunctive Updating and the following claim:

> **Initial Context:** The initial context set C is such that, for each $w \in C_{(1)}$ there is some compositional context that interprets (1) at w that determines a matching dressing intension, and there is some compositional context that interprets (1) at w that determines a mixed dressing intension.

It can be verified that The Disjunctive Account predicts the desired patterns of updating for both Tim Likes Matching and Tim Dislikes Mixed. To illustrate how this works, let's consider how this account predicts the desired patterns of updating for Tim Likes Matching. The same points apply, mutatis mutandis, for Tim Dislikes Mixed.

Claim: The Disjunctive Account predicts that conditions (i) and (ii) of Safe Information are satisfied for Tim Likes Matching.

> To see why this is so, let's consider how the updates generated by this discourse work when we represent this process in terms of the evolution of a set of context-fragments.
>
> Focussing on our context set $C_{(1)}$, there are two relevant ways that the world parameter of a context fragment in this set could be. First, the world parameter could be such that Tim likes matching pairs and dislikes mixed pairs. We'll denote such a world parameter: $w_{\mathrm{Match}^\heartsuit}$. Second, the world parameter could be such that Tim likes mixed pairs and dislikes matching pairs. We'll denote such a world parameter: $w_{\mathrm{Mixed}^\heartsuit}$.
>
> Similarly, there are two relevant ways that the compositional context parameter of a context fragment in the fragmentation of $C_{(1)}$ could be. First, the compositional context parameter could be such that it determines a matching dressing intension. We'll denote such a compositional context parameter: $c_{\mathrm{Match}^{\mathrm{Int}}}$. Second, the compositional context parameter could be such that it determines a mixed dressing intension. We'll denote such a compositional context parameter: $c_{\mathrm{Mixed}^{\mathrm{Int}}}$.
>
> Given these distinctions, we can partition the fragmentation of $C_{(1)}$ into four subsets depending on which of the two types

of world parameters and which of the two types of compositional context parameters the context fragment has. For our purposes, the important point is that, for each of the sentences in our discourse, whether a context fragment is eliminated by the assertion of that sentence in that discourse depends just on which of these four classes that context fragment is in. Given this, we can provide a simplified context fragment model by taking the fragmentation of $C_{(1)}$ to consist of the following four context fragments:

- $\langle c_{\text{Match}^{\text{Int}}}, w_{\text{Match}^{\heartsuit}} \rangle$
- $\langle c_{\text{Mixed}^{\text{Int}}}, w_{\text{Match}^{\heartsuit}} \rangle$
- $\langle c_{\text{Match}^{\text{Int}}}, w_{\text{Mixed}^{\heartsuit}} \rangle$
- $\langle c_{\text{Mixed}^{\text{Int}}}, w_{\text{Mixed}^{\heartsuit}} \rangle$

And we can represent the patterns of updating, given each of our discourses, by representing which of these context fragments is in the relevant updated set of context fragments.

The following table indicates, for each of the context sets that results from each successive assertion in Tim Likes Matching, which of these fragments is in the relevant fragmentation of that set. In each case, one can determine which context fragments remain, given the relevant updates, by consulting the column under the appropriate updated context set.

	$C^{(1)}$	$C^{(1)(2)}$	$C^{(1)(2)(3)}$
$\langle c_{\text{Match}^{\text{Int}}}, w_{\text{Match}^{\heartsuit}} \rangle$	✓	✓	✓
$\langle c_{\text{Mixed}^{\text{Int}}}, w_{\text{Match}^{\heartsuit}} \rangle$	✓	X	X
$\langle c_{\text{Match}^{\text{Int}}}, w_{\text{Mixed}^{\heartsuit}} \rangle$	✓	X	X
$\langle c_{\text{Mixed}^{\text{Int}}}, w_{\text{Mixed}^{\heartsuit}} \rangle$	✓	✓	X

Let's verify that, for each assertion, the updated context set is exactly as desired. First, given the assertion of (1), no context fragments are eliminated, and so no information is imparted. Next, given the assertion of (2), we eliminate $\langle c_{\text{Mixed}^{\text{Int}}}, w_{\text{Match}^{\heartsuit}} \rangle$ and $\langle c_{\text{Match}^{\text{Int}}}, w_{\text{Mixed}^{\heartsuit}} \rangle$. That is, we eliminate those context fragments in which Tim likes mixed pairs of socks, but the compositional context determines a matching dressing intension,

as well as those context fragments in which Tim likes matching pairs of socks but the compositional context determines a mixed dressing intension. This still leaves two context fragments, each witnessing one of Tim's two possible preferences, and so the second assertion also provides no information about his preferences, as desired. Finally, the assertion of (3) eliminates $\langle c_{\mathrm{Mixed}^{\mathrm{Int}}}, w_{\mathrm{Mixed}^{\heartsuit}} \rangle$, leaving only context fragments in which Tim likes matching pairs of socks and dislikes mixed pairs. Thus, we get the desired result that jointly these three assertions communicate the information about Tim's preferences.

The preceding shows that that condition (i) of Safe Information is satisfied for Tim Likes Matching, given The Disjunctive Account. The above reasoning also shows that condition (ii) of Safe Information is satisfied, since every $w \in \mathrm{F}$ in which Tim Likes Matching occurs is the second member of a context fragment of the form: $\langle c_{\mathrm{Match}^{\mathrm{Int}}}, w_{\mathrm{Match}^{\heartsuit}} \rangle$ and so will remain uneliminated by this process of updating.

According to The Disjunctive Account, the distinctive patterns of updating exhibited in Tim Likes Matching are explained by appeal to certain initial metasemantic symmetries present in these discourses together with certain facts about how the assertions of the relevant sentences both eliminate worlds from the context set and eliminate certain ways in which later sentences in these discourses are interpreted. And the same is true for Tim Dislikes Mixed.

The Disjunctive Account provides, I think, an attractive way of predicting the desired patterns of updating in these discourses. Disjunctive Multi-Context Updating and Contextual Pruning provide simple and principled accounts of how context sets are updated and how the interpretation of an assertion in a discourse depends on previous assertions in the discourse. And the assumption codified in Initial Context is one that there is good prima facie reason to endorse. For the Presuppositions would seem to be symmetric with respect to mixed and matching interpretations of (1), and Initial Context respects this apparent metasemantic symmetry.

The fact that Disjunctive Updating allows us to predict in a simple and principled manner the patterns of updating exhibited in Tim Likes Matching and Tim Dislikes Mixed, given an independently plausible story about

which compositional contexts interpret the initial assertion of (1) in these discourses, provides, I think, strong evidence in favor of this account of conversational updating.

Finally, it's worth noting that, in addition to predicting the truth of Safe Information, The Disjunctive Account is also jointly compatible the truth of Preservation, Metasemantic Symmetry and Supervenience, (and so, also, Minimal Symmetry).

This account entails the truth of Preservation, since this follows from Contextual Pruning. Indeed, Contextual Pruning entails the following more general principle:

> **Generalized Preservation:** Given a context set C, and discourse in which α and β are asserted in succession, for each $w \in C_\alpha$ if there is a unique compositional context c that interprets an assertion of α in w, then if $w \in C^\alpha$, then c uniquely interprets the subsequent assertion of β in w.

This account, then, vindicates the plausible thought that, in our discourses, if there is a unique interpretation of a given sentence at a given world in the context set, then this interpretation gets passed along to any subsequent assertion in this discourse at this world of the context set.

This account is also jointly compatible with the truth of Metasemantic Symmetry and Supervenience, since the patterns of interpretation postulated by this account are compatible with (i) it being the case that these facts are determined just by the facts that obtain up to and including the time of the relevant assertion, and (ii) there being worlds in which these discourses occur that are perfectly symmetrical up to and including the time of assertion of their common first sentence.

The present account, then, is not only able to correctly predict the patterns of updating exhibited by our two discourses, but is able to do so in a way that is compatible with our other plausible principles.

4 Comparisons, qualifications and conclusion

In this closing section, I'll say a few things about the relation between Disjunctive Updating and three accounts of updating — proposed in Barker 2002, 2013, MacFarlane 2020 and King 2021 — that bear strong similarities to this account.

As we've seen, Disjunctive Updating can predict the patterns of updating exhibited in the discourses in Sarah's Socks. As we noted above, one of the key features that ensures these predictions is that, given this account, we can view conversational updating as the result of updating a set of context-fragments. Accounts with this same formal structure are also proposed in Barker 2002, 2013 and MacFarlane 2020. These authors should, I think, view the preceding arguments as providing further support for their accounts.

In Barker 2002, 2013, it is suggested that a conversational state may be modelled by appeal to a set of ordered-pairs roughly similar to the ordered pairs we've called *context-fragments*, and that such conversational states may be updated in roughly the manner in which we've taken sets of context-fragments to be updated. In Barker's models, the non-world parameter only settles certain context sensitive matters such as the delineations for gradable adjectives such a 'tall', but I take it that context-fragments provide a natural generalization of this idea.

Barker 2002 is primarily concerned with accounting for the ways in which certain uses of vague predicates such as 'tall' may serve to simultaneously provide information about an individual's height, as well as information about the standards to count as 'tall', and to account for the ways in which such worldly and metalinguistic information may be entangled. Thus, suppose that it is contextually uncertain how tall Owen is — either he is $6'1''$ or $5'11''$. In addition, suppose that it is contextually uncertain what the cut-off for 'tall' is — it may be anywhere from $5'10''$ to $6', 2''$. Given this sort of context, consider an assertion of:

(O) Owen is tall.

As Barker observes, such an assertion will naturally serve to rule out those possibilities in which Owen is $6'1''$ and the cut-off for 'tall' is above $6'1''$, and those possibilities in which Owen is $5'11''$ and the cut-off for 'tall' is above $5'11''$. We can model this by taking the context-set to consist of a set of context-fragments. The assertion of (O) rules out those fragments whose world-parameter determines that Owen is $6'1''$ and whose non-world parameter determines a cut-off for 'tall' above $6'1''$, as well as those fragments whose world-parameter determines that Owen is $5'11''$ and whose non-world parameter determines a cut-off for 'tall' above $5'11''$.

As Barker notes, though, one can in principle model such discourses while allowing that the metalinguistic information is also determined by the world parameter. Barker 2002 is officially neutral on the question of whether, given

an asserted sentence and a world, there is a unique resolution of the relevant context sensitive parameters, and so is neutral on whether, given some set of ordered pairs representing the conversational state, there will be at most one, or potentially more than one, pair $\langle c, w \rangle$, for a given world w. Barker 2013, however, is more sympathetic to latter view. The proponent of this sort of view may then see the data about the patterns of updating exhibited by Tim Likes Matching and Tim Dislikes Mixed as providing further support for their account of conversational updating.

A natural question for this view, though, is how we should interpret such sets of context fragments? What is it for a conversational state to be accurately modelled by appeal to such a set?

A novel way of answering this question has been developed in MacFarlane 2020. There MacFarlane argues that we should model conversational updating by appeal to the sorts of sets of ordered-pairs that Barker proposes. MacFarlane, though, rejects the claim that there is a unique resolution of the relevant contextual parameters given an asserted sentence and a world. According to MacFarlane, then, given some set of ordered pairs representing the conversational state, there may be two ordered pairs that agree with respect to their second world-type parameter, but disagree with respect to their first parameter.

Formally, then, the view proposed in MacFarlane 2020 is very close to the context-fragment model. MacFarlane, though, adopts a certain expressivist interpretation of this formalism. In particular, according to MacFarlane, we should think of the second world-type parameter of such a pair as representing a maximally opinionated belief state, and the first non-world parameter as representing a maximally determinate plan — akin to the more general hyperplans proposed in Gibbard 2003 — that settles the intension of every term. Roughly, then, according to MacFarlane, for a set of context-fragments X to accurately model a conversational state just is for the participants in the conversation to mutually adopt a hybrid cognitive/planning state that may be represented by the set of ordered-pairs in X.

The proponent of Disjunctive Updating provides what is, in principle, a distinct answer to the above question. According to this account, for a set of context-fragments X to accurately model a conversational state just is for X to be the fragmentation of C_ϕ, where C_ϕ is the context-set of the conversation, given an assertion of a sentence ϕ. That is, according to Disjunctive Updating, for a set of context-fragments X to accurately model a conversational state, given an assertion ϕ, just is for it to be the case that X is the set

of pairs $\langle c, w \rangle$ such that it is compatible with the mutual presuppositions of the conversational state, given the assertion of ϕ, that the world is w and the assertion of ϕ at w is *interpreted* by c.

Unlike the account in MacFarlane 2020, this view is not in itself an expressivist theory of updating. To see this, note that Standard Updating is not in itself an expressivist theory of updating. Instead, on a natural interpretation, this account simply postulates a particular word-world relation — *interpretation* — which this account maintains is one-one. This relation, then, plays a particular role in conversational updating. In particular, according to this account, agents have certain attitudes — specifically, instances of the sui generis attitude of presupposition — about the possible extensions of this word-world relation, and the dynamics of conversational updating are understood in terms of such attitudes. This view, in itself, involves no commitment to such a relation having any ultimate analysis in expressivist terms. Similarly, I suggest that Disjunctive Updating, on a natural interpretation, can be seen as appealing to this same word-world relation — though this account maintains that this relation may in principle be one-many — without any commitment to such a relation having any ultimate analysis in expressivist terms.

There is much to be said about MacFarlane's way of thinking about conversational states and how it compares to the account proposed here. A full comparison of these two views, though, must wait for another occasion. It is worth, however, stressing that one thing that the preceding establishes is that, at least insofar as one takes there to be an intelligible non-expressivist understanding of Standard Updating, we do not *need* to think about the sorts of conversational states that arise, given the above sorts of discourses, in expressivist terms. For, given such an understanding of Standard Updating, it seems clear that there is a parallel way of understanding the account developed in the previous section that provides a natural alternative non-expressivist interpretation of the formalism that MacFarlane and Barker both propose. Thus, while one *can* explain the patterns of updating considered in this paper in expressivist terms, one need not.

Finally, King (2021) develops an account of conversational updating that bears important similarities — and important dissimilarities — to the account developed here. Crucially, King 2021 agrees with the present account, and gives a battery of cases that support the claim, that often there are multiple compositional contexts that serve to interpret a given sentence, and so, of-

ten we express multiple propositions in asserting a single sentence. Let me briefly say some things about the relation between these two accounts.

One difference between the account developed here and the account developed in King 2021 is that the latter account is primarily focused on the effect of single assertions on the context set. This account doesn't, then, focus on the way in which certain assertions may systematically constrain the available interpretations for later assertions. In particular, then, there is nothing in King's account that naturally corresponds to Contextual Pruning. As far as I can see, though, such an additional principle could be easily endorsed, and, for the reasons adumbrated above, I'm inclined to think that the patterns exhibited by the discourses we've considered provide good motivation for endorsing this sort of principle.

A more significant difference between the present account and the account developed in King 2021 concerns the way that the context set is updated given a multiplicity of propositions that are expressed. King 2021 agrees with the present account that, in many such cases, we update on the disjunction of the propositions asserted. And, I take it that King would agree with the present account that this is the right picture of how updating proceeds in the cases that we've considered. Importantly, though, King (2021) maintains that in other cases we update on the *conjunction* of propositions asserted. And this, of course, is not predicted by Disjunctive Updating.

Now I think that it's a subtle issue whether the cases that King 2021 discusses — in which it seems that we do update conjunctively — can be handled given a suitably augmented account that endorses Disjunctive Updating. However, an adequate discussion of this must also wait for another occasion. Here I'll simply note two options. The first is to explain the apparent patterns of conjunctive updating by appeal to Disjunctive Updating together with additional metasemantic and pragmatic principles. The second is to take Disjunctive Updating to only hold for a restricted class of cases and to endorse a more general account of updating, perhaps similar to that developed in King 2021, that delivers Disjunctive Updating as the appropriate rule in the relevant restricted class of cases.[13] While I think that it's an important

---

13 King (2021) considers the issue of whether the account proposed in Barker 2013 can adequately predict conjunctive readings. Given the formal similarity between the latter account and Disjunctive Updating, the dialectic is essentially the same when we consider whether the proponent of Disjunctive Updating can adequately predict conjunctive readings. King (2021) concludes, as I do here, that it is unclear whether such an account can be adequately supplemented with additional pragmatic principles to predict the conjunctive readings or not.

question which of these two options is preferable, for present purposes I'll simply note that, in either case, one can see the preceding argument as vindicating Disjunctive Updating, at least for cases similar to the ones that we've considered above.[14]

In summation: I'm inclined to think that Standard Updating has much to recommend it. Consideration of discourses such as Tim Likes Matching and Tim Dislikes Mixed shows, however, that we shouldn't accept this account of conversational updating. In particular, by focussing on configurational predicates such as 'pair of socks' we can see that there are natural discourses whose effects cannot be naturally predicted if one maintains that, at each world in the context set, there is a single interpretation of the relevant predicate. Disjunctive Updating, though, provides a natural amendation of this account that is able to predict the patterns of updating that arise given these discourses. In such cases, we can explain the effects of a given discourse by maintaining that an assertion simultaneously serves to update the shared information in the conversation, as well as to cut down on which compositional contexts are available to interpret later assertions in the discourse.

References

Barker, Chris. 2002. The dynamics of vagueness. *Linguistics and Philosophy* 25(1). 1–36. https://doi.org/10.1023/A:1014346114955.

Barker, Chris. 2013. Negotiating taste. *Inquiry* 56(2-3). 240–257. https://doi.org/10.1080/0020174X.2013.784482.

van den Berg, Martin H. 1996. Discourse grammar and dynamic logic. In Paul Dekker & Martin Stokhof (eds.), *Proceedings of the tenth Amsterdam Colloquium: December 18-21, 1995*, 93–112. Amsterdam: ILLC/ Department of Philosophy, University of Amsterdam.

Braun, David & Theodore Sider. 2007. Vague, so untrue. *Nous* 41(2). 133–156. https://doi.org/10.1111/j.1468-0068.2007.00641.x.

---

14 Another apparent difference between the present account and the account in King 2021 is that, unlike the former, the latter is not diagonal in form — that is, it doesn't consider, for each world in the context set, which propositions are expressed at that world, and, on the basis of this, tell you whether that world should or should not be eliminated from the context set. It is, however, unclear to me to what extent this marks a significant difference between the two accounts. The reason for this is that, for all of the central cases considered in King 2021, it is common ground which compositional contexts interpret a given assertion, and so, at each world in the context set, the same compositional contexts interpret a given assertion. But, given this, there is no need to distinguish between a diagonal and non-diagonal update.

Buchanan, Ray. 2010. A puzzle about meaning and communication. *Nous* 44(2). 340–371. https://doi.org/10.1111/j.1468-0068.2010.00743.x.

Dummett, Michael. 1959. Truth. *Proceedings of the Aristotelian Society* 59(1). 141–162. https://doi.org/10.1093/aristotelian/59.1.141.

Egan, Andy. 2009. Billboards, bombs and shotgun weddings. *Synthese* 166(2). 251–279. https://doi.org/10.1007/s11229-007-9284-4.

von Fintel, Kai. 1994. *Restrictions on quantifier domains*. University of Massachusetts at Amherst dissertation. https://semanticsarchive.net/Archive/jA3N2IwN/fintel-1994-thesis.pdf.

von Fintel, Kai & Anthony Gillies. 2011. 'Might' made right. In Brian Weatherson & Andy Egan (eds.), *Epistemic modals*, 108–130. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199591596.003.0004.

Gibbard, Allan. 2003. *Thinking how to live*. Cambridge, MA: Harvard University Press.

Jackman, Henry. 1999. We live forwards but understand backwards. *Pacific Philosophical Quarterly* 80(2). 157–177. https://doi.org/10.1111/1468-0114.00078.

Kamp, Hans. 1971. Formal properties of 'now'. *Theoria* 37(3). 227–273. https://doi.org/10.1111/j.1755-2567.1971.tb00071.x.

Kaplan, David. 1989. Demonstratives: An essay on the semantics, logic, metaphysics and epistemology of demonstratives and other indexicals. In Joseph Almog, John Perry & Howard Wettstein (eds.), *Themes from Kaplan*, 481–563. Oxford: Oxford University Press.

Kennedy, Christopher. 1999. *The syntax and semantics of gradability and comparison*. New York: Routledge.

Kennedy, Christopher. 2007. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy* 30(1). 1–45. https://doi.org/10.1007/s10988-006-9008-0.

King, Jeffrey C. 2014. The metasemantics of contextual sensitivity. In Alexis Burgess & Brett Sherman (eds.), *Metasemantics*. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199669592.003.0004.

King, Jeffrey C. 2018. Strong contextual felicity and felicitous underspecification. *Philosophy and Phenomenological Research* 97(3). 631–657. https://doi.org/10.1111/phpr.12393.

King, Jeffrey C. 2021. *Felicitous underspecification*. Oxford: Oxford University Press. https://doi.org/10.1111/phpr.12393.

Kirk-Giannini, Cameron Domenico. 2020. Why horizontalism. *Philosophical Studies* 177(10). 2881–2905. https://doi.org/10.1007/s11098-019-01350-9.

Kocurek, Alexander, Ethan Jerzak & Rachel Rudolph. 2020. Against conventional wisdom. *Philosophers' Imprint* 20(22). 1–27. https://doi.org/2027/spo.3521354.0020.022.

Kratzer, Angelika. 2012. An investigation of the lumps of thought. In *Modals and conditionals*, chap. 5, 111–159. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199234684.003.0005.

Krifka, Manfred. 2009. Counting configurations. *Sinn und Bedeutung (SuB)* 13. 309–324. https://doi.org/10.18148/sub/2009.v13i2.552.

Lederman, Harvey. 2014. *Uncommon knowledge*. Oxford University dissertation.

Lewis, David. 1980. Index, context, and content. In Stig Kanger & Sven Ohman (eds.), *Philosophy and grammar*, 79–100. Dordrecht: D. Reidel. https://doi.org/10.1007/978-94-009-9012-8_6.

MacFarlane, John. 2020. Indeterminacy as indecision. *Journal of Philosophy* 117(11). 593–664. https://doi.org/10.5840/jphil202011711/1240.

Montague, Richard. 1968. Pragmatics. In Raymond Klibansky (ed.), *Contemporary philosophy: A survey, volume 1*, 102–122. Florence: La Nuova Italia Editrice.

Muñoz, Patrick. 2020. Behavioral attitudes in hyperintensional grammar. University of Chicago manuscript.

Ninan, Dilip. 2010. Semantics and the objects of assertion. *Linguistics and Philosophy* 33(5). 355–380. https://doi.org/10.1007/s10988-011-9084-7.

Rabern, Brian. 2013. Monsters in Kaplan's logic of demonstratives. *Philosophical Studies* 164(2). 393–404. https://doi.org/10.1007/s11098-012-9855-1.

Stalnaker, Robert. 1978. Assertion. In Peter Cole (ed.), *Pragmatics* (Syntax and Semantics 9), 315–332. New York: Academic Press. https://doi.org/10.1163/9789004368873_013.

Stalnaker, Robert. 1996. On the representation of context. *Semantics and Linguistic Theory (SALT)* 6. 279–294. https://doi.org/10.3765/salt.v6i0.2768.

Stalnaker, Robert. 2004. Assertion revisited: On the interpretation of two-dimensional modal semantics. *Philosophical Studies* 118(1-2). 299–322. https://doi.org/10.1023/b:phil.0000019550.81145.34.

Stalnaker, Robert. 2014. *Context*. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199645169.001.0001.

Stanley, Jason. 1997. Rigidity and content. In Richard G. Heck (ed.), *Language, truth, and logic.* Oxford: Oxford University Press. https://doi.org/10.1093/oso/9780198239208.003.0005.

Stanley, Jason & Zoltán Gendler Szabó. 2000. On quantifier domain restriction. *Mind & Language* 15(2-3). 219–261. https://doi.org/10.1111/1468-0017.00130.

von Stechow, Arnim. 1984. Comparing semantic theories of comparison. *Journal of Semantics* 3(1-2). 1–77. https://doi.org/10.1093/jos/3.1-2.1.

Stonić, Una. 2017. On the connection between semantic content and the objects of assertion. *Philosophical Topics* 45(2). 163–180. https://doi.org/10.5840/philtopics201745218.

Vlach, Frank. 1973. *'Now' and 'then': A formal study in the logic of tense anaphora.* University of California Los Angeles dissertation. https://philpapers.org/archive/VLANAT.pdf.

Yalcin, Seth. 2007. Epistemic modals. *Mind* 116(464). 983–1026. https://doi.org/10.1093/mind/fzm983.

Michael Caie
170 St. George St
Toronto, Ontario, M5R 2M8, Canada
caiemike@gmail.com