CONTEXTOLOGY

Simon Goldstein

Cameron Domenico Kirk-Giannini

[Preprint — Please cite published version.]

Abstract

Contextology is the science of the dynamics of the conversational context. Contextology formulates laws governing how the shared information states of interlocutors evolve in response to assertion. More precisely, the contextologist attempts to construct a function which, when provided with just a conversation's pre-update context and the content of an assertion, delivers that conversation's post-update context. Most contextologists have assumed that the function governing the evolution of the context is simple: the post-update context is just the pre-update context intersected with the content of the assertion. We argue that this assumption is wrong: not only is it false, it is also incoherent given standard contextological assumptions. Moreover, it is impossible in principle to revise it to correctly describe the dynamics of context. We conclude that there can be no science of Contextology. The laws governing the evolution of the context in response to assertion must make essential reference to the private information states of interlocutors.

1 Introduction

Many theories of linguistic communication appeal to the notion of *context*. These theories model communicative exchanges using the evolution of a formal object (the context) representing the information shared by speakers and hearers. The body of information which constitutes the context of a conversation at a time is derived somehow from the mental states of the participants of the conversation. When someone in the conversation makes an assertion, the context changes to incorporate the informational content of her assertion.

Notable work which appeals to a context in this sense includes the pragmatic theory of Stalnaker (1973, 1974, 1978, 1998, 2002, 2014, 2018, etc.), as well as Karttunen (1974), Gazdar (1981), Heim (1992), Barker (2002), Egan (2007, 2018), Roberts (2012), Murray (2014), Yalcin (2012, 2018), Bowker (2019), and Cariani (fc), among others. For critical discussion of the notion of context, see Harris (2019).[1]

We formalize the central commitments of this picture as follows:

(1) **(Reduction)** The context c of a conversation C is a function of the private information states of the participants of C.

(2) **(Public Functionality)** The context c' of a conversation C immediately after an assertion with content p is accepted is a function of p and the context c of C immediately before the assertion is accepted.[2]

Assertions do not only change states of shared information; if they are to be of any practical utility to individuals, they must also change the private information states of conversational participants. It is therefore natural to accept:

(3) **(Private Functionality)** The private information state i' of a participant S in conversation C immediately after she accepts an assertion with content p is a function of p and S's information state i immediately before she accepts the assertion.[3]

[1] Another tradition of work, stemming from Kaplan (1977/1989) and Lewis (1981), appeals to a distinct notion of context. Instead of using the word 'context' to refer to bodies of information, this tradition uses it to pick out the features of concrete utterance situations which settle the values of various context-sensitive expressions. Our arguments below do not engage with this alternative conception of context.

[2] Note that we use 'accept' here to describe the act a speaker performs with respect to a given assertion just in case she does not reject it. The act of rejecting an assertion is understood by contextologists to be a move in the conversational language game which any interlocutor can perform in response to an assertion and which prevents that assertion from updating the context. Thus Stalnaker (2014): "...it is a rule of the assertion game that the addressee has the option of rejecting the assertion, blocking this rule-governed change [to the context]. Rejection is another of the possible rule-governed moves in the game." Acceptance in this sense should be carefully distinguished from the propositional attitude of acceptance which Stalnaker (2014) uses to define the notion of a conversational context. We use the word 'accept' in both ways in what follows; in each case, our intended meaning should be clear from context.

[3] It is worth noting that, while theorists who posit a conversational context often explicitly discuss and endorse (Reduction) and (Public Functionality), (Private Functionality) is usually taken for granted rather than explicitly discussed. The idea expressed in (Private Functionality) — that the way a rational agent's attitudes evolve in response to new information is determined by her existing information together with the content of the new information — is nevertheless ubiquitous: one encounters it in a wide range of formal

(1–3) enforce that the context of a conversation and the private information states of its participants evolve together while respecting the determination relation between the former and the latter. According to this picture, there are lawlike generalizations about how the context of a conversation will evolve in response to an assertion, given only information about the content asserted and the context immediately before the assertion is accepted. In other words, this is a picture on which there is a genuine *science* of the evolution of the context in response to assertion. We refer to this science of the dynamics of the conversational context as **Contextology**.

The laws of Contextology describe the *essential dynamic effect* of assertion. That is, they model the evolution of information states in response to assertion *after* taking into account what interlocutors learn simply by observing the speaker's assertoric utterance (the "manifest event" of the assertion). For the act of asserting a proposition, like any other publicly observable act, results in interlocutors learning that the speaker has uttered a certain sentence, spoken with a particular accent, and so forth. This sort of learning from an assertion occurs whether it is ultimately accepted or rejected by interlocutors. The laws of Contextology, in contrast, describe the way the content of an assertion changes public and private information states when it is accepted. When we refer to the information state *immediately before an assertion is accepted*, we mean the state which incorporates whatever information interlocutors learn from the manifest event of that assertion, but which has not yet been updated with its content.[4]

(1–3) collectively imply a further constraint: that two different ways of changing the context always deliver the same result. The first way directly applies (Public Functionality):

theories of belief revision, whether qualitative (as in, for example, Alchourrón, Gärdenfors, and Makinson (1985)) or probabilistic (as in standard Bayesian approaches). For this reason, we believe that all or nearly all theorists who posit a conversational context would endorse (Private Functionality). Thanks to an anonymous referee for pressing us to clarify this point.

[4]Robert Stalnaker explicates the notion of the information state immediately before an assertion is accepted as follows: "In general, given the assumption that the semantics of the language and the rules of the game are common knowledge among the players, we can conclude that when an assertion is made, it will be a manifest event that an assertion with a particular content was made by the person who made it... The context in which an assertion is evaluated (the point at which the addressee decides whether to accept or reject it) is the context as it is after the assertion has been made and mutually recognized." (2014, p. 51). We discuss the significance of the information interlocutors gain from observing the manifest events of assertoric utterances in section 5.

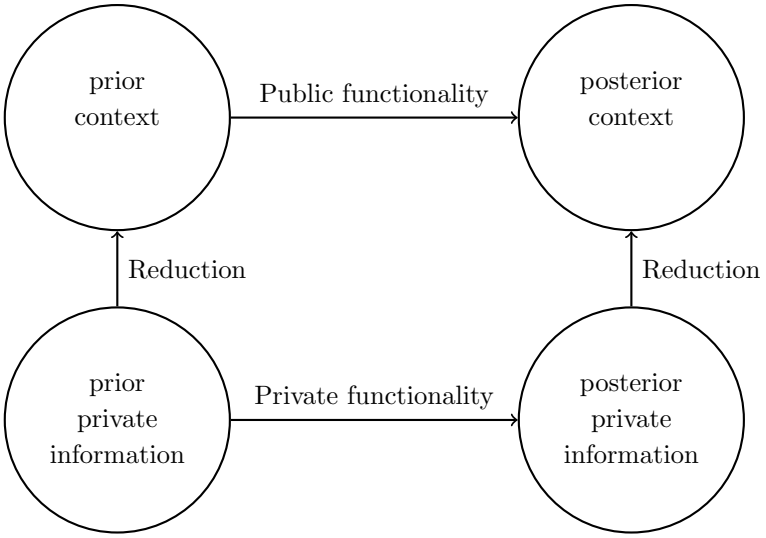

Figure 1: Coherence

the new context after an assertion of p is a function of the old context and p. The second way of changing the context applies (Private Functionality) and then (Reduction). To find the new context, start by applying (Private Functionality) and updating each individual's information state with p. Second, apply (Reduction) and create a new context from the new information of each individual. We can state this crucial corollary as follows:

(4) **(Coherence)** The context c' of a conversation C immediately after an assertion with content p is accepted is identical to the context determined by the private information states of every participant S in C immediately after accepting that assertion.[5]

(Reduction) is not specific concerning the nature of the function which takes us from the information states of interlocutors to the contexts of their conversations. Similarly, (Public Functionality) and (Private Functionality) are not specific concerning the nature of the functions which govern the evolution of public and private information states. So there are many theories of communication which satisfy all three principles. Nevertheless,

[5]More precisely: let c be the context, and let I be the set of information states of the interlocutors of the context. Let $\sqcup$ be the function introduced by (Reduction), which maps interlocutors' information states to contexts. Let $+$ be the function introduced by (Private Functionality), which maps private information states and contents to new private information states. Let $\cdot[\cdot]$ be the function introduced by (Public Functionality), which maps contexts and contents to new contexts. Then (Coherence) says that $c[p] = \sqcup\{i + p \mid i \in I\}$.

we believe it is possible to establish two surprising generalizations: First, many of the most widely accepted contextological theories violate (Coherence); second, (Public Functionality) is untenable. We refer to the first of these problems as *the problem of incoherence*, and the second as *the foundational problem*.

We present models demonstrating the problem of incoherence and the foundational problem below. To foreshadow, the lesson we propose to draw from these problems is that the public information state of a group does not in general encode enough information to determine what information the group will have after updating on the content of an assertion. For this reason, there is no way to describe how the context of a conversation will evolve in response to an assertion of without describing in detail the private information states of the interlocutors in that conversation. A proper theory of communication should not try to construct a function describing how the context evolves in response to assertion, because there isn't one. Rather, a proper theory of communication should give a systematic characterization of how the private information of conversational participants changes in response to successful assertions.

Our discussion proceeds as follows. Section 2 summarizes the most influential version of Contextology. Sections 3 and 4 introduce the problem of incoherence and the foundational problem. Our presentation of the two problems in these sections is simplified; in the following two sections, we consider whether introducing additional complexity affords the contextologist a way to solve them. Section 5 addresses complications which arise when one considers the changes in interlocutors' private information states resulting from the observation that a speaker has asserted. Section 6 addresses complications which arise when care is taken to ensure that the result of updating a public or private information state in response to assertion is always another information state. Section 7 concludes by considering the significance of our arguments for dynamic theories of meaning.

A word of warning. Many of our arguments in what follows do not strictly rely on just the three principles above. Rather, throughout we show that the various natural ways of fleshing out these principles that have been suggested so far in the literature are unsustainable. The

foundational problem in section 3 comes closest to objecting to these principles without further assumptions, but even that argument has a few small background assumptions.

2 The Possibility-Carving Vision

Contextology is dominated by a research paradigm we call *The Possibility-Carving Vision*, developed most systematically by Robert Stalnaker.[6] The Possibility-Carving Vision is distinguished by its account of what makes (Reduction) and (Public Functionality) true.

2.1 Reduction

We begin with (Reduction). Here, The Possiblity-Carving Vision says that the conversational context is defined in terms of the *common* version of whatever propositional attitude $\mathcal{A}$ is used to characterize interlocutors' private information states. That is, possibility carvers assume that there is some propositional attitude $\mathcal{A}$ such that the context of a conversation consists of just those propositions p such that: all interlocutors bear $\mathcal{A}$ to p, all interlocutors bear $\mathcal{A}$ to the proposition that all interlocutors bear $\mathcal{A}$ to p, and so forth:

(5) **(Common Reduction)** There is some propositional attitude $\mathcal{A}$ such that, for any conversation C among interlocutors $i_1...i_n$, the context of C consists of all and only the propositions which are common-$\mathcal{A}$ among $i_1...i_n$.

Throughout, we accept (Common Reduction).[7] Our arguments will not turn on which attitude exactly $\mathcal{A}$ is. For simplicity, we often talk as if $\mathcal{A}$ is belief. In that case, the context is the set of worlds consistent with the interlocutors' common beliefs.[8]

[6] See, for example, Stalnaker (1973, 1974, 1978, 1998, 2002, 2014, 2018).

[7] Lederman (2018a) and Lederman (2018b) argue that common attitudes are elusive, and that various phenomena explained by common attitudes can be explained in other ways. Here, it is worth noting that our arguments below can be adapted straightforwardly to apply to other iterated notions of context. For example, consider the proposal that the context is the set of worlds consistent with what the interlocutors all believe that the interlocutors all believe (Lederman (2018b, p. 1095) discusses proposals of this type). Our arguments below can be adapted straightforwardly to apply to this proposal, and others like it.

[8] See for example Stalnaker (2002, p. 704). Stalnaker (2014) models context in terms of common acceptance rather than common belief. Finally, Stalnaker (2002, p. 716) defines the context as whatever is accepted to be commonly believed. Our arguments below also apply to this more complex conception of context, since we can simply assume in our examples that each agent believes a claim if and only if she accepts it. Because

In accepting The Possibility-Carving Vision's conception of (Reduction), we choose not to consider certain alternative ways of thinking about the context. For example, we do not consider the proposal that the context is identical with the private information state of the hearer or the proposal that the context is the set of propositions believed by all interlocutors (their *mutual beliefs*).

We are not troubled by this choice for two reasons. First, these alternative ways of understanding the context have not been as influential as The Possibility-Carving Vision's conception. Second, these alternative ways of understanding the context seem to make incorrect predictions about when certain speech acts are felicitous. There are a number of plausible principles linking the public information state of a conversation with the set of speech acts it is felicitous for interlocutors to perform. For example, if it is public information that p, it is not felicitous for interlocutors to inquire whether p or to assert either p or $\neg p$. Similarly, if it is not public information that p, it is generally infelicitous to produce an utterance which presupposes p (barring accommodation). But even if both speaker and hearer believe p, it can make sense for speaker to assert p if she believes that hearer does not believe p. This suggests that the set of mutual beliefs of all interlocutors is not a good candidate for the public information state of a conversation. Similarly, identifying the conversational context with the hearer's beliefs makes the unpalatable prediction that it is infelicitous to assert p if one disagrees with one's interlocutor about whether p. These predictive failures are not trivial: making correct predictions about felicity is a central component of the theoretical role the notion of the public information state of a conversation is introduced to play. Alternative conceptions of the context have been less influential than the Possibility-Carving Vision precisely because of this kind of predictive inadequacy.

Possibility carvers also typically assume that $\mathcal{A}$ satisfies various logical principles. First, they require positive and negative introspection. That is, they assume that any individual who bears $\mathcal{A}$ to p must also bear $\mathcal{A}$ to the proposition that she bears $\mathcal{A}$ to p, and any individual who does not bear $\mathcal{A}$ to p must also bear $\mathcal{A}$ to the proposition that she does not

our results in what follows are negative, we only need one such context in order to make our argument.

bear $\mathcal{A}$ to p.[9] Second, they assume that $\mathcal{A}$ is consistent, so that any agent who bears $\mathcal{A}$ to p must fail to bear $\mathcal{A}$ to $\neg p$, and that $\mathcal{A}$ is closed under logical consequence. Below, only part of our discussion (the material about failures of preservation in section 6) hinges on these particular assumptions. Our main arguments can be run whether or not $\mathcal{A}$ is introspective. But for simplicity, throughout we only introduce examples in which these conditions are met.[10]

Throughout, it will be helpful to model both individual and common attitudes using modal logic. We follow The Possibility-Carving Vision in holding that both states of information and the contents of assertions can be modeled as sets of possible worlds. Consider a conversation C taking place between two agents, Speaker and Hearer, at a world w. We can model the context of C in terms of the iterated beliefs of Speaker and Hearer. So let W be a set of possible worlds. Following the textbook semantics (Hintikka (1962)), let S and H be accessibility relations on W that reflect what Speaker and Hearer believe: Speaker believes p at w just in case p is true at v whenever wSv; Hearer believes p at w just in case p is true at v whenever wHv. We let Sw and Hw denote the speaker and hearer's possibilities at w.

The context is the set of worlds consistent with what Speaker and Hearer commonly believe (what they both believe, and both believe they believe, and so forth). Common belief corresponds to the transitive closure R of Speaker and Hearer's accessibility relations. That is, wRv iff there is some chain of worlds $x_1, \ldots x_n$ where $wXx_1 \ldots$ and x_nXv, and X is either S or H in each instance. Speaker and Hearer commonly believe p at w just in case p is true at v whenever wRv.

Here is an informal way of thinking about the context which we find helpful: the context of Speaker and Hearer's conversation at a world w is the set of worlds reachable by starting

[9] For example, Stalnaker writes, "I assume that one has access to what one accepts [for the purposes of the conversation], which means that one accepts that one accepts that ψ when one accepts it, and accepts that one does not when one does not" (2014, p. 45). For arguments that these assumptions are too strong in this setting, see Hawthorne and Magidor (2009).

[10] Stalnaker often restricts his attention to contexts that satisfy an extra constraint. These "non-defective" contexts are ones where communication proceeds especially smoothly. Stalnaker (2004) defines the notion of non-defectiveness for contexts such that a context is non-defective just in case every agent in the context believes it is commonly believed that p iff it actually is commonly believed that p. In other words, non-defective contexts are those where no agents are mistaken about the context. Throughout, we confine our attention to non-defective contexts.

at w and crossing some finite number of "bridges" (accessibility arrows from either S or H). For the purposes of defining the context, it doesn't matter whether the bridges crossed belong to Speaker's accessibility relation or Hearer's. This is because the context models what *both Speaker and Hearer* can take for granted for the purposes of the conversation. For example, even if a world is not accessible from w according to Speaker, Speaker cannot take for granted that that world is not accessible for the purposes of the conversation if it is accessible according to Hearer from any world accessible from w according to Speaker. After all, for all Speaker believes in such a situation, Hearer regards w as a genuine possibility. That w is not a genuine possibility is part of Speaker's private beliefs, but not part of the context.[11]

2.2 Public Functionality

We now turn to (Public Functionality). The Possibility-Carving Vision holds that the evolution of the context in response to assertions is governed by *simple intersection*:

(6) **(Naive Public Update)** If c is the context of a conversation immediately before an assertion with content p is accepted, and if the assertion is accepted, then the context c' immediately after the assertion is accepted is $c \cap p$.

Here are two famous affirmations of this principle:

> To make an assertion is to reduce the context set in a particular way, provided that there are no objections from the other participants in the conversation. The particular way in which the context set is reduced is that all of the possible situations incompatible with what is said are eliminated. (Stalnaker 1978, p. 86)[12]

[11] Thanks to an anonymous referee for pressing us to clarify this point.

[12] Similarly, here is Stalnaker: "Assume that the speech act of assertion is governed by the following rule: an assertion changes the context by adding the propositional content of the assertion to the common ground" (2014, 51). And here is Yalcin: "The effect of successful assertion, on Stalnaker's view, is to eliminate possibilities incompatible with the proposition asserted from the common ground" (2012, p. 2).

> "Meaning determines the content of an assertion as a function of context, and the assertion rule takes the prior context set to a posterior context set, which is the intersection of the prior set with that content." (Stalnaker 1993, p. 11)[13]

We argue below that (Naive Public Update) is unacceptable. But it will help to appreciate what makes The Possibility-Carving Vision so attractive if we pause first to reflect on its virtues. First, it is uncontroversial that the communicative effect of an accepted assertion results at least in that assertion's content being included in the interlocutors' public information state. (Naive Public Update) is the most direct possible explanation of this fact. It holds that the public information state is updated with the content of the assertion and nothing else.

Second, if a speaker asserts p and this assertion is accepted by all conversational participants, subsequent assertions of p or $\neg p$ become infelicitous because redundant (p) or inconsistent with what is taken for granted ($\neg p$). Combined with our earlier observations about the ways in which the set of felicitous continuations of a conversation is constrained by its context, (Naive Public Update) explains this fact because it requires that the context following the acceptance of an assertion of p must include p.

In attributing to The Possibility-Carving Vision a commitment to (Public Functionality) and (Naive Public Update), we do not suggest that its advocates deny the platitudinous observation that the act of asserting a proposition can in various senses communicate more than the content of that assertion. First, as mentioned above, there is the information revealed to a speaker's interlocutors by the manifest event of her assertoric utterance. Second, narrowing the context in response to an assertion of p can cause propositions other than p to be commonly believed if they are contextually entailed by p, where proposition q is contextually entailed by proposition p in context C just in case q is true at every world in C where p is true. This sort of communication does involve the essential dynamic effect of assertion but does not call into question either (Public Functionality) or (Naive Public

[13]See also Heim (1992, p. 185) and Murray (2014, p. 22), who follow Stalnaker in holding explicitly that assertions are proposals to update the common ground according to (Naive Public Update).

Update). Our cases have a different structure. Our claim is that (Public Functionality) and (Naive Public Update) fail because there are cases in which the context of a conversation after an assertion of p comes to entail a proposition other than p, but not because that proposition is contextually entailed by p in the context prior to updating on the assertion. *This* sort of failure has not been appreciated by friends of The Possibility-Carving Vision or in the broader literature on assertion.[14]

It is worth noting that our arguments in what follows leave a number of important aspects of The Possibility-Carving Vision intact. Nothing we argue conflicts with there being a conversational context which governs the felicity of assertions of sentences with presuppositions, for example. The mistake is to suppose that the context's evolution in response to assertion can be described without essential reference to the private propositional attitudes of interlocutors.

3 The Problem of Incoherence

In this section and the next, we introduce two problems for Contextology. We keep our presentation of the problems in these sections as simple as possible; readers who suspect that we have *over*simplified matters are invited to look ahead to the discussion in sections 5 and 6.

We have seen that The Possibility-Carving Vision endorses the simplest possible public update rule which guarantees that accepting an assertion of p results in a context where p is true. The corresponding private update rule is:[15]

[14]Two other examples: First, asserting a sentence which carries a presupposition can, via a process of accommodation, cause that proposition to become commonly believed. Second, asserting a sentence containing a conventional implicature item like a non-restrictive relative clause can result in the implicated proposition becoming commonly believed. These processes do not involve the essential dynamic effect of assertion, and our arguments do not make reference to them.

[15]Like (Private Functionality), (Naive Private Update) is a principle which is more often assumed than explicitly endorsed by contextologists. It is, however, simple and independently plausible — it corresponds exactly to the *Expansion* rule in the AGM theory of belief revision (Alchourrón, Gärdenfors, and Makinson (1985)), for example — and we believe many contextologists endorse it. Note, however, that our arguments succeed even on alternative conceptions of private update; see section 6 for discussion. Thanks to an anonymous referee for pressing us to clarify this point.

(7) **(Naive Private Update)** If i is the private information state of an interlocutor immediately before an assertion with content p is accepted, and if the assertion is accepted, then that interlocutor's information state i' immediately after the assertion is accepted is $i \cap p$.[16]

In this section, we show that combining (Naive Public Update) with (Naive Private Update) results in violations of (Coherence). This is the problem of incoherence.

Consider the conversation between Speaker and Hearer represented in Figure 2.[17,18]

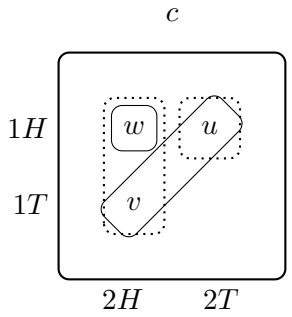

Figure 2: Incoherence

The thickly outlined regions represent doxastic accessibility for Speaker, while the dotted regions represent doxastic accessibility for Hearer. In this example, accessibility is an equivalence relation: each world in a region "sees" all and only the worlds within that region. Suppose that the actual world is w. Then Speaker believes $1H$ and $2H$, while Hearer

[16]Strictly speaking, we model an agent's information at world w with an accessibility relation, not simply a set of worlds. We can derive (Naive Private Update) from a corresponding update procedure on accessibility relations. Suppose that in world w and time t an interlocutor's belief state is represented by the accessibility relation S (so that her belief state at w at that time is Sw). Suppose that between t and t' the interlocutor learns that p. Then the agent's information at t' can be represented by the accessibility relation $S' = S \cap \{\langle w', v\rangle \mid w' \neq w \text{ or } v \in p\}$.

[17]This example is adapted from van Ditmarsch, van der Hoek, and Kooi (2007, p. 82). They use this example to demonstrate a principle that is structurally related to the failure of (Coherence): that the formula $[p]Cq \leftrightarrow (p \rightarrow C[p]q)$ is invalid in Public Announcement Logic. Here, the left-hand side roughly says that q is accepted in the state of common knowledge that results from each agent learning from a public announcement of p, and the right-hand side roughly says that if p then all p worlds consistent with what is commonly known are q worlds. To our knowledge, however, no one in the Public Announcement Logic tradition has appreciated the implications of examples of this kind for formal pragmatics. We regard it as unfortunate that there has not been more dialogue between philosophers working on the foundations of conversation, on the one hand, and logicians in the DEL tradition, on the other. Indeed, we see our arguments against Contextology as attempts to integrate the insights of these two theoretical approaches.

[18]For simplicity, we omit from the diagram the world (z) where both propositions are false, since it is not in the transitive closure of the epistemic accessibility relations of Speaker and Hearer at w.

believes $2H$ but is agnostic about $1H$. The worlds consistent with what Speaker and Hearer commonly believe at w are $\{w, v, u\}$.

Here is a case which models context c.

> [Coins]: Two fair coins are flipped. If both coins land tails, a loud alarm goes off, which both Speaker and Hearer would hear. A tiny bell rings in Speaker's ear just in case both coins land heads. A tiny bell rings in Hearer's ear just in case the second coin lands heads. Speaker and Hearer commonly believe that this is the setup. Both coins actually land heads.

Let $1H$ be the claim that the first coin lands heads and $2H$ be the claim that the second coin lands heads ($1T$ and $2T$ are the corresponding claims about landing tails). After both coins have been flipped and landed heads, the only public information is that the loud alarm didn't go off, indicating that at least one coin has landed heads. But both Speaker and Hearer also have private information. Speaker's personal bell rings, informing her that both coins have landed heads — this is why Speaker's doxastic accessibility relation has word w seeing only itself. Hearer's personal bell also rings, informing him that the second coin has landed heads. He remains ignorant, however, concerning whether the first coin landed heads: his doxastic possibilities at w are w and v. Crucially, since Hearer is not sure whether the first coin landed heads, he is not sure whether Speaker's personal bell rang. But he knows that if it did not ring, then Speaker's information (no loud alarm; no personal bell) allows her to conclude that exactly one coin landed heads, but not which. This is why Speaker's doxastic possibilities at world v include only worlds v and u. Thus, though both Speaker and Hearer can privately rule out one or both of worlds v and u, the context, which models what can be taken for granted for the purpose of the conversation, consists of all three worlds at which one or more coins landed heads.

Now suppose Speaker asserts $1H$ and Hearer accepts her assertion. Since c consists of just worlds w, v, and u, the result of applying (Naive Public Update) to it is $\{w, u\}$. Evolving the private beliefs of Speaker and Hearer using (Naive Private Update), on the

other hand, will have the following effects: Speaker's beliefs about the coins will not change, since Speaker already believes $1H$. Hearer, on the other hand, will rule out world v, since v is a $1T$-world. Moreover, since the setup of the case is commonly believed between Speaker and Hearer, Speaker believes that Hearer will come to believe $1H\&2H$ upon ruling out world v, Hearer believes that if $1H\&2H$, then Speaker's ear-bell must have gone off, so that Speaker believes $1H\&2H$, and so forth. It follows that after the update world w is the only world accessible for both Speaker and Hearer: neither Speaker nor Hearer considers it possible that either coin landed tails, or that the other considers it possible that either coin landed tails, and so forth. This means that it is commonly believed in the state resulting from Speaker's assertion that $1H\&2H$. It follows that the result of applying (Naive Public Update) to c is not the same as the result of evolving the private information states of Speaker and Hearer using (Naive Private Update) and then pooling the result. Naive updating violates (Coherence).

[Coins] and cases like it have the following structure: Some worlds where $1H$ is true are only consistent with what is commonly believed because they are doxastic possibilities at some $1T$ world. If we once again think of the set of worlds in Speaker and Hearer's public information state at w as the set of worlds reachable by starting at w and crossing some finite number of "bridges," this observation corresponds to the idea that some $1H$-worlds are included in the context only because of a "bridge" which runs through a $1T$-world. When Speaker and Hearer update their private information states with $1H$, all $1T$-worlds are eliminated, and these bridges are destroyed. This is how updating on $1H$ can come to rule out some $1H$-worlds in addition to all $1T$-worlds.

4 The Foundational Problem

In this section, we build on the problem of incoherence, arguing that it is not possible to study the evolution of the context as the contextologist tries to. One lesson of [Coins] is that the context of a conversation is not rich enough to represent dependency relationships

of the sort that explain how updating on $1H$ can come to rule out some $1H$-worlds. Given the context of a conversation, there is no way to determine whether some $1H$-worlds are compatible with what is commonly believed only because some $1T$-worlds are compatible with what is commonly believed. This suggests that it is not possible to predict the evolution of what is commonly believed in response to new information based only on a complete description of the context and the information learned.

Our argument has the following form. We introduce two different conversations with the same context and argue that updating on the content of an assertion will cause the private information states of the interlocutors in the two conversations to evolve in different ways, so that the resulting sets of common beliefs are not identical. If this argument is successful, it demonstrates that the evolution of the context in response to an assertion is not a function of what is commonly believed before the assertion is accepted together with the content asserted. In other words, (Public Functionality) is false, and Contextology is impossible. This is the foundational problem.[19]

To begin, consider the two conversations in Figure 3.

The first conversation is the [Coins] example from earlier. The second conversation is a variant of [Coins]. In this variant, a tiny bell rings in Speaker's ear just in case the first coin lands heads, but Hearer has no special information.

Speaker and Hearer have different beliefs in these two conversations. But they have the same common beliefs: in each case, they commonly believe just that at least one coin landed heads. The context in both conversations is $\{w, v, u\}$.

But now consider how each context would change in response to Speaker asserting that

[19]Our argument holds fixed (Common Reduction) and (Naive Private Update). One could in principle seek to salvage Contextology by constructing the context out of interlocutors' private information states in a less direct way, or by endorsing a different private update rule, or by abandoning (Reduction) or (Private Functionality) altogether. We find these options unattractive: As long as the context plays its usual theoretical role of representing the information interlocutors take for granted for the purposes of the conversation, it seems to us that in both of our models below, the context must contain at least the worlds we specify prior to being updated and exactly the worlds we specify after being updated — so doing away with (Common Reduction) is no help. Along similar lines, we argue in section 6 that there is no plausible private update rule on which our argument would not succeed. Thus, although (Common Reduction) and (Naive Private Update) are officially premises of our argument, we think the foundational problem arises for a wide variety of possible views about the nature of the context and the dynamics of private information states — though not, perhaps, for *all* possible views.

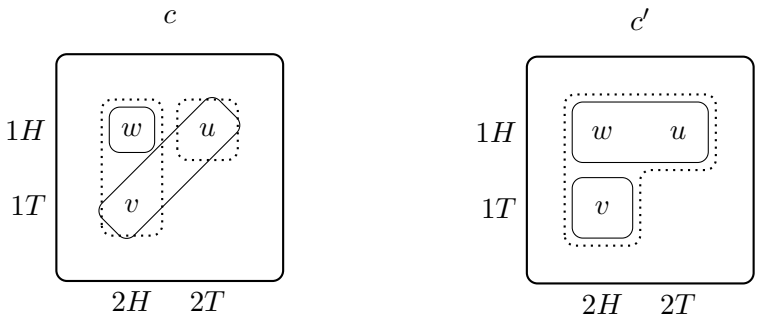

Figure 3: Same context, different beliefs

the first coin landed heads. Above, we saw that (Naive Private Update) predicts that in the first conversation, Hearer will respond to an assertion of $1H$ by entering the state $\{w\}$. After this update, Speaker and Hearer's new common beliefs would be modeled by $\{w\}$. By contrast, consider the result of applying (Naive Private Update) to the second conversation. In this case, Hearer's initial belief state is merely $\{w, v, u\}$, and so after updating with $1H$ he would enter state $\{w, u\}$. In the second conversation, Speaker's belief state before and after updating on $1H$ would be $\{w, u\}$. So after their respective private updates, Speaker and Hearer's new common beliefs would be modeled by $\{w, u\}$. Thus, each conversation produces a different context after Speaker and Hearer privately update on $1H$, at least according to (Naive Private Update).

We believe that the problem of incoherence and the foundational problem constitute decisive objections to Contextology. Devoted contextologists may, however, object at this point that our simplified presentation of the arguments has ignored certain important aspects of the way the context evolves in response to assertion, and has somewhat tendentiously presupposed the simplest possible public and private update rules for assertion. These are reasonable concerns. As we argue in the following sections, however, taking them into account does not diminish the force of our objections to Contextology.

5 Manifest Update

As we have seen, the public and private information states updated in response to an assertion are not the same as the public and private information states of interlocutors before the assertion occurs. The manifest event of an assertion changes public and private information states in a way which may affect interpretation. For example, the manifest event of an assertion might make it clear that the speaker presupposes something, thereby causing the audience to presuppose it, thereby rendering felicitous an otherwise defective utterance. In this section, we consider whether accounting for updating on manifest events of assertion resolves the problem of incoherence and the foundational problem.

What exactly do interlocutors learn from observing the manifest event of an utterance? We assume they learn at least enough for the resulting context to entail that the speaker believes the content asserted. In other words, we assume it is common-$\mathcal{A}$ after the utterance that the speaker believes the content asserted (and therefore, by positive introspection, Ω-believes it).[20] Moreover, it is plausible that each interlocutor s learns two material conditionals: first, that if s accepts the speaker's assertion of p, then s will believe (and Ω-believe) p; second, that if all interlocutors accept the speaker's assertion of p, then p will be included in the conversation's public information state. These conditionals will become important in the next section.

It might at first seem that appealing to manifest update offers the contextologist an escape from the problem of incoherence. Return to [Coins]. The manifest event of Speaker asserting that the first coin landed heads allows Hearer to infer that Speaker believes that the first coin landed heads. Hearer believes that Speaker would have this belief only if it were true (since Hearer believes that Speaker will remain agnostic about which coin landed heads unless her personal bell rings, in which case she will know that both coins landed heads). So observing the manifest event of Speaker's utterance puts Hearer in a position to infer that the first coin, and hence both coins, landed heads. It follows that [Coins] is a case where the manifest event of Speaker's assertion is enough to narrow the context to

[20] An agent Ω-believes p iff they believe p, believe that they believe p, and so forth.

just world w. So applying naive intersective update rules to the context and to Speaker and Hearer's private information states will have no effect, and the case is not a counterexample to (Coherence).

But it is possible to modify [Coins] to generate a counterexample. For example, suppose Hearer is unsure of how Speaker will respond to the coin flips. Hearer entertains two hypotheses. The first hypothesis is as before: that a tiny bell rings in Speaker's ear iff both coins land heads, and that Speaker believes this and will respond reasonably to silence by suspending judgment concerning which coin has landed heads. Hearer's second hypothesis is that Speaker will not respond to silence in the correct way. If only one coin lands heads, Speaker will immediately infer that it is the first coin. Both the setup and Hearer's suspicions about Speaker are common knowledge between the two. As before, both coins lands heads. Figure 4 summarizes the relevant structure.

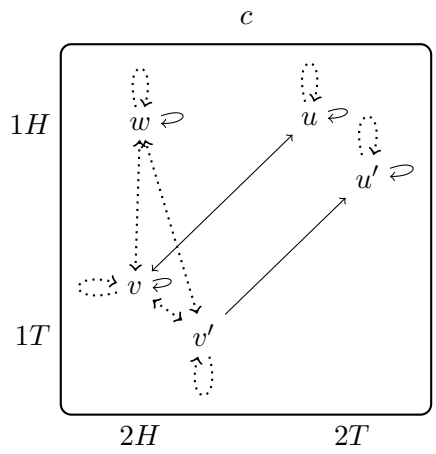

Figure 4: Before the manifest update

The thick lines represent Speaker's doxastic accessibility relation, while the dotted lines represent Hearer's. Prime worlds differ from their unmarked counterparts only in Speaker's beliefs. At both v and v', the first coin lands tails and the second lands heads. In v, Speaker correctly infers that exactly one coin has landed heads, and nothing more. In v', Speaker incorrectly infers that the first coin landed heads and the second landed tails. Thus at world w, where both coins land heads and Speaker does not make any mistakes, Hearer believes

it is possible that the first coin landed tails and the second coin landed heads, and also that Speaker made an error in that case, believing that the first coin landed heads and the second landed tails. The context of the conversation at world w thus includes five worlds: $\{w, v, v', u, u'\}$.

Updating on the manifest event of Speaker's assertion that the first coin landed heads does not eliminate all $1T$ worlds. In particular, since Hearer only learns that Speaker believes that the first coin landed heads, Hearer does not rule out world v'. But in v', the first coin landed tails, even though Speaker believes it landed heads. So the context after updating on the manifest event consists of w, v', and u', represented in Figure 5.

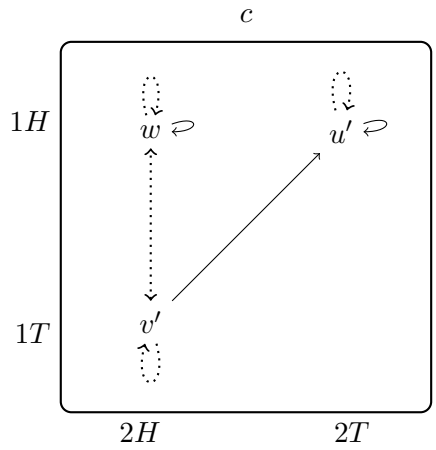

Figure 5: After the manifest update

Applying (Naive Public Update) to this context eliminates just the one $1T$ world v'. However, applying (Naive Private Update) and then pooling the result eliminates both v' and u'. To see this, note that Speaker's information at w is already the singleton of w. When Hearer updates with $1H$, he rules out world v'; so w is the unique world consistent with what Speaker and Hearer now commonly believe. Once again we have a violation of (Coherence).[21]

[21] Is this modified version of [Coins] suspicious in some way? For example, we have stipulated that Hearer accepts Speaker's assertion even though he leaves open the possibility that Speaker has not responded correctly to silence. But perhaps there is some general cooperative principle governing conversation which requires us to take for granted (at least for the purposes of conversation) that our fellow interlocutors are responding correctly to their evidence. Or perhaps there is some general cooperative principle governing

Given that appealing to manifest update does not solve the problem of incoherence, it should come as no surprise that it does not solve the foundational problem, either. To see why, contrast the situation depicted in Figure 4 with one in which Hearer entertains the same two hypotheses about Speaker but has no special information about the coins. In this modified case, manifest update on Speaker's assertion leaves worlds w, v' and u' accessible for Hearer at w. Thus the context after manifest update is the same as in Figure 5, but when Hearer updates his private information state with $1H$, he is not able to rule out world u' (because he does not have access to his ear-bell). So we have the same post-manifest-update context in the two cases, but private update on the content asserted has different effects.

conversation which secures that updating on an assertion results in it becoming commonly believed that the speaker who produced it was responding correctly to their evidence. The first of these suggestions seeks to dismiss our modified version of [Coins] on the basis that the context as we have described it fails to meet some constraint on what sort of contexts are rationally permissible even before Speaker asserts. The second seeks to dismiss our modified versions of [Coins] on the basis that in cases of rational communication the context will not evolve as we have suggested.

Three points are relevant here. First, nothing about the structure of the case requires that Hearer leaves open the possibility that Speaker does not respond to silence in the correct way. For example, instead of imagining that Hearer leaves open the possibility that Speaker will respond incorrectly to silence, we can imagine that Hearer believes it is possible that Speaker has yet another tiny bell: one wired to ring just in case the first coin comes up heads. However, hearer also believes that if Speaker has this bell, and if the first coin lands tails and the second lands heads, a once-in-a-lifetime intervention by the demon Chort will cause the bell to ring in Speaker's ear, giving him misleading evidence that the first coin has landed heads and the second coin has landed tails. If Speaker has this extra bell and the first coin comes up heads, on the other hand, its ringing gives Speaker knowledge that the first coin has landed heads and the second tails. In this case, Hearer is sure that Speaker would respond to silence in the correct way, though he is not sure that Speaker will assert truly. And the revised case can be used to demonstrate a violation of (Coherence) in the same way as the original.

What is crucial to the case, then, is the fact that Hearer leaves open the possibility that Speaker asserts *falsely* even if she is responding correctly to her evidence. Here the second point comes into play: If we require all interlocutors to rule out from the get-go the possibility that speakers might assert falsely, we trivialize the theory of assertion. For then learning that a speaker has asserted p suffices for learning p itself, and the essential effect of assertion is irrelevant to the dynamics of the context. (Even setting this trivialization worry aside, it strikes us as implausible that conversational participants do or should rule out the possibility that their interlocutors will assert falsely.)

For these reasons, we do not think the first strategy of dismissing [Coins] by placing rational constraints on what the context can be like before Speaker asserts is promising. What of the second strategy of building certain extra requirements into the theory of updating? Dialectically, this second strategy is premature: our argument in this section is just that accounting for manifest update does not solve the problem of incoherence for *naive* update rules. But more importantly (and this is the third point), adding some stipulation to the effect that updating on Speaker's assertion must involve coming to believe that she responded correctly to her evidence will not solve the problem, since (as we have seen) we can imagine that at world u' she asserts reasonably and even knowingly. This means that even the imagined strengthened public update rule will not eliminate world u' from the context.

Still other variants of [Coins] might involve agents who fail to infer anything about another speaker's evidence from observation of the manifest event of their assertion — thanks to an anonymous referee for emphasizing this point.

The foundational problem persists.

Summarizing, we've seen that the problem of incoherence and the foundational problem also affect theories of context that take into account the changes brought about by the manifest events of assertoric utterances.

6 Advanced Update

Manifest update does not solve the problem of incoherence or the foundational problem. In this section, we consider whether more sophisticated update rules might solve the problem instead. Indeed, we will see that there are good reasons (what we call *failures of preservation*) to replace naive updating with more advanced update rules. However, we will show that the problem of incoherence and the foundational problem persist for a variety of more complex update rules and argue that in fact they must persist for any acceptable update rule.

Here is why naive intersective update rules need to be replaced anyways. The structure of common $\mathcal{A}$ and the contextologist's principles of positive and negative introspection for private information states place constraints on what bodies of information could possibly count as the context of a conversation or the $\mathcal{A}$-attitudes of an interlocutor. Both (Naive Private Update) and (Naive Public Update) fail to guarantee that the result of updating a private or public information state in response to an assertion is an information state at all. We call this sort of situation a *failure of preservation*; failures of preservation are a serious challenge to simple update rules of the sort which are ubiquitous in the literature on formal pragmatics.[22]

With respect to (Naive Private Update), failures of preservation arise as follows: Suppose we have an individual with coherent private information state s who is agnostic about p and also about whether she will ever believe p, in such a way that her private information state contains at least one world where p is true and she never comes to believe p. If we hold that

[22]When it comes to private information states, cases involving failures of preservation are structurally similar to the paradoxes of introspective belief change familiar from dynamic epistemic logic. See Enqvist and Olsson (2013) for a helpful overview. Our own discussion of these cases differs from extant work in connecting these issues to the theory of context. We discuss the relevance of work in dynamic epistemic logic to our arguments later in this section.

private information states are updated by intersection, we have the consequence that upon updating on an assertion of p, our individual moves to belief state $s \cap p$. But intersecting s with p does nothing to change the fact that our individual is agnostic about whether she will ever believe p. In particular, intersecting s with p does not rule out the worlds in s where p is true but she never comes to believe p. So while p is true throughout $s \cap p$ (and thus our individual believes p), it is not true throughout $s \cap p$ that she believes p (and so she fails to believe that she believes p). In other words, $s \cap p$ fails to satisfy the principle of positive introspection, and is therefore not a private information state at all.

With respect to (Naive Public Update), note that because the conversational context is defined in terms of common $\mathcal{A}$, the commitments of the context satisfy positive introspection. In other words, where c is the context of a conversation, and Cp represents that p is commonly believed by the interlocutors of that conversation, we have:

(8) **(Context Introspection)** If p is true at every world in c, then Cp is true at every world in c.[23]

But the result of updating a context with the content of an assertion via (Naive Public Update) will not necessarily result in an information state which is introspective. This is because (Naive Public Update) does not eliminate the possibility that c contains a world where the content of the assertion is true but is never commonly believed. Indeed, this will usually be the case: with respect to most propositions p which are not entailed by the conversational context, it is natural for interlocutors to believe that p could be true while Cp is false. This is just to admit the banal fact that the truth of p does not by itself guarantee that p becomes common information: it must be believed by at least one interlocutor, this interlocutor must assert it, the assertion must be accepted, and so forth. But if it is common information that if p, it might or might not ever come to be true that

[23]To see why common belief satisfies positive introspection, recall that the common beliefs of two agents at w are found by taking the transitive closure R of the accessibility relations S and H of the two agents at w. Here the key observation is that R is itself a transitive relation. If some proposition p is true throughout the set of worlds Rw accessible from w, then it must also be true throughout the set of worlds Rw' accessible from any world w' in Rw, since Rw' must be a subset of Rw.

Cp, the context contains a world where p and $\neg Cp$, and updating it in accordance with (Naive Public Update) will fail to eliminate this world, resulting in a state which violates (Context Introspection). Such a state cannot be the result of pooling the $\mathcal{A}$ attitudes of any set of interlocutors; it is not even possibly a conversational context. Thus neither public nor private naive update rules form a coherent theoretical package when combined with the introspection principles endorsed by The Possibility-Carving Vision.

A natural response to this observation is to revise both rules so that updating a private information state in response to an assertion of p involves learning that one has accepted the speaker's assertion that p, and updating a public information state in response to an assertion of p involves learning that the assertion has not been rejected. This yields the following update rules:

(9) **(Introspective Private Update)** If i is the private information state of interlocutor a immediately before an assertion with content p is accepted, and if a accepts this assertion, then a's information state i' immediately after the assertion is accepted is $i \cap [p \,\&\, (a \text{ has accepted the speaker's assertion that } p)]$.[24]

(10) **(Introspective Public Update)** If c is the context of a conversation immediately before an assertion P with content p is accepted, and if the assertion is accepted, then the context c' immediately after the assertion is accepted is $c \cap [p \,\&\, (P \text{ has not been rejected})]$.

Since we have assumed that interlocutors learn from the manifest event of an assertion that they will believe its content if they choose to accept it, and that this content will become commonly believed if the assertion is not rejected, (Introspective Private Update) and (Introspective Public Update) ensure that updating on the content of an assertion does not lead to failures of preservation.

[24] In the presence of our assumptions about manifest update, (Introspective Private Update) is at least as strong as the alternative update rule which requires i to update her private information state with p and the proposition that she Ω-believes p. For discussion of a similar rule, see Segerberg (2006). See Gerbrandy and Groeneveld (1997), van Benthem (2007), and Caie (2019) for other accounts of updating which imply that whenever an agent updates on p, they subsequently Ω-believe p.

The optimistic contextologist might hope at this point that moving to introspective update rules will solve the problem of incoherence or the foundational problem. Unfortunately, this is not the case. Consider again the context in our revised version of [Coins] after updating on the manifest event of Speaker's assertion that the first coin landed heads, represented in Figure 5.[25] Evolving the public and private information states in this case using (Introspective Private Update) and (Introspective Public Update) rather than (Naive Private Update) and (Naive Public Update) does not change the result: (Naive Private Update) was already enough to get each private information state to the singleton set of world w, and the result of evolving the context using (Introspective Public Update) is the same as the result of evolving it using (Naive Public Update), since what comes to be entailed in the context when it is updated with the information that Speaker's assertion has not been rejected is just that it is commonly believed that p, and this was true after evolving the context using (Naive Public Update), as well. For the same reason, introspective update rules fail to solve the foundational problem, as well.

The problem of incoherence and the foundational problem also arise with other advanced update rules. Rather than relying on our assumption that updating on the manifest event of an assertion causes all interlocutors to believe that if that assertion is not rejected, its content will come to be commonly believed, one might prefer an alternative update rule which directly updates c with both p and the proposition that it is commonly believed that p:

(11) **(Communal Public Update)** If c is the context of a conversation immediately before an assertion with content p is accepted, and if the assertion is accepted, then the context c' immediately after the assertion is accepted is $c \cap (p \,\&\, Cp)$.

But note that, if our private update rule does not enforce that interlocutors update on a proposition at least as strong as Cp, we once again run the risk of violating (Coherence)

[25] We focus on the revised version of [Coins] in this section because, in light of the discussion in Section 5, it is more probative than the original version of [Coins]. But it is worth noting that the advanced update rules presented in this section do not escape the problems with the original version of [Coins], either.

because our public update rule eliminates more worlds from the context than pooling the results of our private update rules. So (Communal Public Update) is naturally paired with:

(12) **(Communal Private Update)** If i is the private information state of an interlocutor immediately before an assertion with content p is accepted, and if the assertion is accepted, then that interlocutor's information state i' immediately after the assertion is accepted is $i \cap (p \,\&\, Cp)$.

Might the combination of (Communal Public Update) and (Communal Private Update) do what naive and introspective update rules could not and solve the problem of incoherence and the foundational problem? No: Our modified [Coins] case in Figure 5 still provides a counterexample to (Coherence) on the assumption that public and private information states evolve in accordance with communal update rules. The problem for (Coherence) is that Speaker and Hearer learn *more* from Speaker's assertion than that it is commonly believed that p. So enforcing that they learn that it is commonly believed that p could not in principle fix the problem.

Might there perhaps be some further update rules which rescue Contextology from counterexample? We think not, because any plausible story about private updating will deliver the same anti-contextological results. First, suppose that updating is weakly eliminative, so that each agent's new belief state is a subset of their earlier state whenever there is a subset of the earlier state that is consistent with the proposition being learned. It is then uncontroversial that Hearer's posterior belief state in the first conversation in Figure 3 will be $\{w\}$. If any fewer worlds were contained in the state, Hearer's beliefs would be absurd; if any more worlds were contained in the state, Hearer would not learn anything. Similarly, it is uncontroversial that Hearer's posterior belief state in the second conversation will be $\{w, u\}$: otherwise, he would arbitrarily form a belief about $2H$ when learning about $1H$, despite having no information that correlates the two claims.[26] It follows that no defensible private update rule will validate (Coherence) or solve the foundational problem. If this is

[26] All of the private update rules we have canvassed deliver these results.

correct, Contextology is a mistake. A single context set can produce two different contexts after updating with $1H$, depending on the private information states of interlocutors.[27]

Even more complex theories of updating deliver similar results in this case. Dynamic epistemic logic (DEL) models a similar kind of case, where agents update their beliefs in response to a 'public announcement' of p.[28] To model such cases, DEL represents the beliefs of the various agents with a model containing a set of possibilities, accessibility relations over those possibilities, and an assignment of truth values to various sentences in a language. DEL represents public announcements in terms of changes to this model. When Speaker announces p, DEL updates the model representing the beliefs of Speaker and Hearer. The worlds in the updated model are just those in the original where p is true. Crucially, the new accessibility relations are the restriction of the earlier accessibility relations to the worlds where p is true. In the first conversation before updating on $1H$, we have for example wHv and vHv. In the model after updating on $1H$, we still have wHv but we do not have vHv: v is not included in this model. It turns out that this procedure agrees exactly with our pretheoretic judgments, delivering $\{w\}$ and $\{w, u\}$ as the posterior contexts of the two conversations in Figure 3.[29]

7 Conclusion

One open question concerns the extent to which our arguments spell trouble for dynamic theories of meaning. For example, Heim (1992) and Veltman (1996) model the semantic effect of presupposition and epistemic modals in terms of changes to various bodies of information. In both cases, they suggest that non-modal information updates a body of information in accordance with the naive update rule considered above.

[27] For simplicity, we frame this argument in terms of the simple examples in Figure 3, but analogous points can be made about the more complex examples in section 5.

[28] For an overview of dynamic epistemic logic, see van Ditmarsch, van der Hoek, and Kooi (2007) (see chapter 4 for an introduction to Public Announcement Logic). For initial applications to public announcements, see Plaza (1989); for more complex development, see for example Baltag et al. (1998).

[29] DEL avoids our results because it has no commitment to (Public Functionality). Rather, DEL explains how each conversational interlocutor's beliefs change in response to a public announcement; the evolution of the context is parasitic on this underlying change.

Here, care is required. In personal correspondence, Veltman suggests that his semantics for epistemic modals can be understood as providing rules for *individual* agents to update their beliefs. No appeal to a common body of information is necessary.[30] By contrast, Heim's theory of presupposition seems to depend on updating a common body of information. In Heim's theory, the meaning of a sentence is a partial function from contexts to contexts: p presupposes q when this function is only defined for contexts that imply q.[31] If we interpret this theory in terms of updating purely private bodies of information, we risk reaching the wrong results about presupposition. For reasons similar to those which led us to reject the proposal that the conversational context is the mutual beliefs of interlocutors, whether it is appropriate to assert a sentence that presupposes q must depend on more than just the speaker and hearer each privately believing q. In order for the assertion to be felicitous, it is important that the two agents commonly believe (or accept) that q. For this reason, we suggest that our arguments above may raise trouble for extant dynamic theories of presupposition, although they may leave dynamic theories of epistemic modals untouched.

Though we do not ultimately endorse any particular rule for how individuals update their private information in response to assertions, we believe that any adequate theory of updating should treat the two cases in section 4 differently. In one case, there is an essential dependence between learning $1H$ and learning $2H$; in the other case, there is not. This is enough to show that there is no science of Contextology in the sense in which we have been using the term.

If Contextology must be abandoned, how should we study linguistic communication? The lesson of the failure of (Public Functionality) is that the theory of assertion must focus on the private information states of interlocutors. We find (Introspective Private Update) and perhaps also (Communal Private Update) to be plausible theories of the dynamics of private information states. Alternatively, one of the many theories of private updating in

[30] This interpretation is also suggested at various points in the text of "Defaults in Update Semantics" (1996). For example, Veltman writes: "I want the information states σ to represent... what the agent regards as his or her knowledge" (p. 260).

[31] For further elaboration of a dynamic treatment of presupposition which faces related problems, see Beaver (2001).

dynamic epistemic logic cited above might have the right story. A satisfactory theory of information exchange via assertion might start with one of these rules and show how to use it to predict how the context of an arbitrary conversation will evolve in response to an assertion given the content asserted and the private information states of its interlocutors. This is not to say that no theoretical role remains for the context in formal pragmatics. As we remarked in section 2, none of our arguments bear on the idea that a conversation's public information state determines which propositions can be felicitously presupposed, which assertions are informative, and so on. For this reason, it remains possible to imagine a successor to Contextology — a *Neo-Contextology*, one might say — which continues to recognize the importance of the context for certain theoretical purposes but appeals to the dynamics of private information states in explaining how the context evolves in response to assertion.[32]

Acknowledgments We are grateful for helpful feedback from Bob Beddor, Kyle Blumberg, Dan Hoek, Ben Holguin, Harvey Lederman, and Bob Stalnaker, from the audience at the Philosophy of Language Digital Work-in-Progress Group organized in December 2020 by Lukas Lewerentz and Henry Schiller, and from two anonymous referees for *Philosophical Studies*.

References

Alchourrón, C.E., Gärdenfors, P. and Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic* 50: 510–530.

Baltag, Alexandru, Moss, Lawrence S., and Solecki, Slawomir. (1998). The logic of common knowledge, public announcements, and private suspicions. In Gilboa, Itzhak (ed.), *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge* (TARK 98), pp. 43–56.

[32]Thanks to an anonymous referee for emphasizing the possibility of Neo-Contextology.

Barker, Chris. (2002). The dynamics of vagueness. *Linguistics and Philosophy* 25: 1–36.

Beaver, David I. (2001). *Presupposition and Assertion in Dynamic Semantics.* CSLI Publications.

van Benthem, Johan. (2007). Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics* 17(2):129–155.

Bowker, Mark. (2019). Saying a bundle: Meaning, intention, and underdetermination. *Synthese* 196(10): 4229–4252.

Caie, Michael. (2019). Doxastic logic. In Pettigrew, Richard and Weisberg, Jonathan (eds), *The Open Handbook of Formal Epistemology.* PhilPapers.

Cariani, Fabrizio. (fc). Assertion and Modality. In Goldberg, Sanford (ed.), *The Oxford Handbook of Assertion.* Oxford University Press.

van Ditmarsch, Hans, van der Hoek, Wiebe, and Kooi, Barteld. (2005). *Dynamic Epistemic Logic.* Springer.

Egan, Andy. (2007). Epistemic modals, relativism and assertion. *Philosophical Studies* 133:1–22.

———. (2018). De Se Pragmatics. *Philosophical Perspectives* 32 (1):144–164.

Enqvist, Sebastian, and Olsson, Erik J. (2013). Segerberg on the Paradoxes of Introspective Belief Change. In Trypuz, Robert (ed.), *Krister Segerberg on Logic of Actions.* Springer.

Fagin, Ronald, Halpern, Joseph Y., Moses, Yoram, and Vardi, Moshe Y. (1995). *Reasoning About Knowledge.* MIT Press.

Gazdar, Gerald. (1981). Speech act assignment. In Joshi, Aravind K., Webber, Bonnie L., and Sag, Ivan A. (eds), *Elements of Discourse Understanding*, Cambridge University Press.

Gerbrandy, Jelle and Groeneveld, Willem. (1997). Reasoning about information change. *Journal of Logic, Language, and Information* 6:147–169.

Harris, Daniel W. (fc). We talk to people, not contexts. *Philosophical Studies.*

Hawthorne, John, and Magidor, Ofra. (2009). Assertion, context, and epistemic accessibility. *Mind*, 118, 377–397.

Heim, Irene. (1992). Presupposition projection and the semantics of attitude verbs. *Journal of Semantics* 9, 183–221.

Hintikka, Jaakko. (1962). *Knowledge and Belief*. Cornell University Press.

Kaplan, David. (1989) Demonstratives. In Almog, Joseph, Perry, John, and Wettstein, Howard (eds), *Themes from Kaplan*, Oxford University Press, pp. 481–564.

Karttunen, Lauri. (1974). Presupposition and linguistic context. *Theoretical Linguistics* 1:181–194.

Russell, Jeffrey S., Hawthorne, John, and Buchak, Lara. (2014). Groupthink. *Philosophical Studies* 172(5): 1287–1309.

Lederman, Harvey. (2018a) Two paradoxes of common knowledge. *Noûs* 52(4): 921-945.

———. (2018b). Uncommon knowledge. *Mind* 127(508): 1069-1105.

Lewis, David. (1981). Index, context and content. In Kanger, Stig and Öhman, Sven (eds), *Philosophy and Grammar*, D. Reidel Publishing Company, pp. 79–100.

Murray, Sarah E. (2014). Varieties of update. *Semantics & Pragmatics* 7(2): 1–53.

Plaza, Jan A. (1989). Logics of public communications. In Emrich, M.L., Pfeifer, M.S., Hadzikadic, M., and Ras, Z.W. (eds), *Proceedings of the Fourth International Symposium on Methodologies for Intelligent Systems*, pp. 201–216.

Segerberg, Krister. (2006). Moore problems in full dynamic doxastic logic. In Malinowski, Jacek, and Pietruszczak, Andrzej (eds), *Poznan Studies in the Philosophy of the Sciences and the Humanities: Essays in Logic and Ontology*, Rodopi, pp. 95–110.

Stalnaker, Robert. (1973). Presuppositions. *Journal of Philosophical Logic*, 2:447–457.

———. (1974). Pragmatic Presuppositions. In his *Context and Content*, Oxford University Press, 1999, pp. 47–62. Originally published in Munitz, Milton K. and Unger, Peter (eds), *Semantics and Philosophy*, New York University Press.

———. (1978). Assertion. In his *Context and Content*, Oxford University Press, 1999, pp. 78–95. Originally published in *Syntax and Semantics* 9: 315-332.

———. (1998). On the representation of context. *Journal of Logic, Language, and Information*, 7:3–19.

———. (2002) Common ground. *Linguistics and Philosophy*, 25:701–721.

———. (2014). *Context.* Oxford University Press.

———. (2018). Dynamic Pragmatics, Static Semantics. In Fogal, Daniel, Harris, Daniel W., and Moss, Matt (eds) *New Work on Speech Acts*, Oxford University Press, pp. 384–399.

Veltman, Frank. (1996) Defaults in update semantics. *Journal of Philosophical Logic*, 25(3): 221–261.

Yalcin, Seth. (2012). Context probabilism. *Lecture Notes in Computer Science* 7218: 12–21.

———. (2018). Expressivism by force. In Fogal, Daniel, Harris, Daniel W., and Moss, Matt (eds), *New Work on Speech Acts*, Oxford University Press, pp. 400–429.