

Epistemic Bubbles and Contextual Discordance

JOEY POLLOCK

Abstract

Recent work in social epistemology has drawn attention to various problematic social epistemic phenomena that are common within online networks. Nguyen (2020) argues that it is important to distinguish epistemic bubbles from echo chambers. An epistemic bubble is an information structure that merely lacks information or sources that would be relevant or important to the user. An echo chamber is a structure in which dissenting opinions are, not necessarily absent, but actively undermined, for example by instilling attitudes of distrust towards their adherents. Because of this, echo chambers are thought to be especially difficult to escape. In contrast, according to Nguyen, it is relatively easy to shatter an epistemic bubble: one simply introduces the missing information. In this paper, I argue that it is more difficult to shatter an epistemic bubble than has been recognised in the literature. The reason for this is the relationship between epistemic bubbles and interpretative resources. Despite their epistemic drawbacks, it is comparatively easy to gain knowledge from sources inside one's epistemic bubble because agents within a bubble share common ground. In contrast, it can be very difficult to gain knowledge from sources outside of one's bubble because interlocutors on the outside are less likely to have the shared context needed to facilitate communicative success. I argue that this problem suggests a different way to understand the nature of epistemic bubbles and our prospects for escaping them.

1. Introduction

Social media is not just a forum for keeping in touch with friends and family, but is now also a primary source of news for many people; as such, it exerts a massive influence on the information that its users consume. Most social media applications are, however, ultimately designed to maximise profit, rather than to promote epistemic goods. Recent work in social epistemology has drawn attention to various problematic social epistemic phenomena that are common within online networks. For example, agents interacting through social media may find themselves in so-called 'epistemic bubbles' due to a combination of their own circumstances, choices, and the algorithms

doi:10.1017/S0031819124000093 © The Author(s), 2024. Published by Cambridge University Press on behalf of The Royal Institute of Philosophy. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

that control what content is most visible to them. Epistemic bubbles are problematic because they lack information or sources that would be relevant or important to the user. Nguyen (2020) distinguishes epistemic bubbles from a related, more insidious, phenomenon: the ‘echo chamber’. An echo chamber is an environment in which dissenting opinions are, not necessarily absent, but actively undermined, for example by instilling attitudes of distrust towards their adherents. These phenomena can serve to stifle the spread of knowledge to those who need it and may contribute to the phenomenon of knowledge resistance, where information is available to an agent, yet not accepted.

Social epistemologists have begun to develop strategies to counteract the problematic effects of social media networks (e.g., Habgood-Coote, [forthcoming](#); Record and Miller, 2022; Ashton, 2020; Frost-Arnold, 2021). In this paper, I draw on work in philosophy of language and pragmatics to investigate a communicative barrier to these attempts to epistemically improve our online environments: contextual discordance. It is well established that an audience may fail to recover an agent’s intended message if there are disparities in the interlocutors’ perceptions of the conversational context (e.g., Bach, 2012; Carston, 2009; Bezuidenhout, 2002). In the literature, some authors have begun to explore the distinctive ways in which these disparities can undermine communication in online contexts (Record and Miller, 2022; Frost-Arnold, 2021). This paper continues this trend by arguing that contextual discordance in online communication presents a special challenge when it comes to escaping epistemic bubbles. Despite their epistemic drawbacks, it is relatively easy to gain knowledge from sources within one’s epistemic bubble, where there is knowledge to be had. This is because, by the very nature of an epistemic bubble, one often shares many contextual assumptions with those inside it. In contrast, it can be difficult to gain knowledge from sources outside of a bubble because interlocutors on the outside are less likely to have the shared interpretative and speech dispositions needed to facilitate communicative success. I argue that this problem suggests a different way to understand the nature of epistemic bubbles and our prospects for escaping them.

2. Nguyen on Epistemic Bubbles and Echo Chambers

The term ‘epistemic bubble’ is used to refer to a number of related phenomena in the literature (see Sunstein, 2018; Pariser, 2011). For Nguyen (2020), an epistemic bubble is a social epistemic structure that merely omits relevant information, or sources thereof, from an

Epistemic Bubbles and Contextual Discordance

agent's network. Bubbles are formed, not necessarily by a deliberate attempt to limit or exclude information, but simply because agents exert some degree of control over who they connect with on social media; and these people – typically friends, family, colleagues, or those with whom the agent shares some common interest – will in many cases have similar beliefs and values to the agent's own, at least within particular domains of discourse. Of course, it is not always the case that we have much in common with family, friends, and colleagues, but those in our social media networks with whom we disagree can be easily 'muted' to hide their content from our feeds, leaving only posts from sources that we approve of, or agree with. Choosing who we connect with online is perhaps not in itself problematic. However, epistemic problems arise when our online social networks are also used as significant sources of information about the world. For many of us, much of the news we are exposed to comes in the form of online articles and posts that have been shared by people we have chosen to populate our social networks. The composition of these networks may differ across social media platforms. For example, on Facebook, our networks will likely be dominated by our friends, family, and co-workers; on platforms like Twitter and Instagram, by contrast, our networks may include many more members whom we do not know personally (either online or offline). Where we choose who populates our social media feeds, the result is a self-imposed filter on the nature of the information we consume, which tends to merely confirm, rather than challenge or extend, our existing worldview. As Nguyen writes, 'we impose on ourselves a narrowed and self-reinforcing epistemic filter, which leaves out contrary views and illegitimately inflates our epistemic self-confidence.' (Nguyen, 2020, p. 142)

Agents may exert a significant degree of control over what information they consume offline too, of course. For example, we decide which newspapers to read, and we can often avoid talking in person to the people we disagree with, or about topics that we disagree over. Thus, social media's impact on the development of our epistemic bubbles is partly a matter of degree. However, there are also important differences in kind when it comes to the manner in which information is selected for our consumption: on many social media platforms, the epistemic filter that we choose to impose on ourselves is amplified by algorithms that feed us whatever content best succeeds in keeping us from disengaging.¹ Moreover, users are

¹ It is worth noting that algorithms do not necessarily feed us content that is consistent with our beliefs or values in order to keep us engaged.

typically unaware of how these algorithms work (Nguyen, 2020, p. 144; Record and Miller, 2013).

Nguyen argues that it is important to distinguish epistemic bubbles from echo chambers. Echo chambers, on Nguyen's view, are social epistemic structures in which contrary information or dissenting sources are actively discredited. Following Jamieson and Cappella (2008), Nguyen claims that the shared beliefs of those within an echo chamber include reasons to distrust those on the outside (2020, p. 146). Echo chambers thus serve to prevent uptake of certain kinds of information. Because of this, unlike an epistemic bubble, information that conflicts with an agent's worldview may in fact exist within an echo chamber. The problem is that this information is rejected: it is ignored, devalued, or undermined, and may even play a role in strengthening the ideology prevalent within the chamber, and the conviction of its members.

Nguyen argues that these two phenomena have typically been lumped together in the literature but that it is vital to distinguish between them because the kinds of harms they cause are distinct and the ways to counteract their problematic effects will differ (p. 142).² Within an epistemic bubble, the problem is mere lack of exposure to information – epistemic bubbles lack 'coverage reliability' (Goldberg, 2011): they expose the agent only to information that is deemed important and interesting by others in the network, but this does not necessarily include all information that would be important or relevant for the agent herself. As such, Nguyen suggests that introducing the missing information – and thus improving coverage reliability – may be enough to burst a bubble. In contrast, exposure to new information, on its own, would not enable an agent to escape an echo chamber as the echo chamber encourages the agent trapped inside to discredit information from outside sources. Given this understanding of the differences between these structures, Nguyen (2020, p. 154) claims that escaping an epistemic bubble is comparatively straightforward. In contrast, escaping an echo chamber is far more difficult: dissenting voices will reinforce an echo chamber, and so mere exposure to information may be, at best, useless and, at worst, counterproductive.

In this paper, my primary aim is to argue that it is more difficult to escape an epistemic bubble than has been recognised in the literature. The reason for this is the relationship between epistemic bubbles, on the one hand, and the conditions conducive to successful

² As Nguyen (2020, p. 142) notes, this distinction is a conceptual one: a social network can exhibit both structures to varying degrees.

Epistemic Bubbles and Contextual Discordance

communication, on the other. Roughly, the idea is this: an agent's interpretative resources cannot outstrip the information that is available to her within her epistemic bubble. Interpreting others within the same bubble is relatively easy because interlocutors within a bubble will typically possess similar interpretative dispositions (at least with respect to certain domains). However, interpreting sources from outside a bubble can be very difficult because the agent is thereby attempting to communicate with an interlocutor in the absence of shared contextual information or common ground. As I will argue, mere exposure to information will not help to burst a bubble because the agent may still lack sufficient resources to build a shared context for interpretation. To show this, I will first explain a range of challenges common in online communication that stem from differences in the interlocutors' respective contextual assumptions. I will then argue for a particular view of the relationship between these communicative issues and the epistemic problems described by Nguyen. It is worth noting that I don't intend my argument in what follows as a sharp criticism of Nguyen's work. I take it that his point in emphasising the relative ease with which we can escape an epistemic bubble is to highlight the difficulty of escaping an echo chamber in comparison, and the distinctive nature of the challenge involved. Nonetheless, I think it is important to attend to obstacles to escaping epistemic bubbles that have been neglected in the literature. I will ultimately argue that epistemic bubbles are not best thought of as structures that we can burst or shatter at all. Rather, epistemic bubbles mark the horizon of our epistemic perspective on the world: they are structures that we can expand, but never escape.

3. Communication and Communicating Online

Whether online or offline, communication involves substantial reliance on context. A context of utterance is the environment or conversational setting in which the utterance is produced. It includes both external physical features of this environment – perhaps just those that are epistemically accessible – as well as facts about the interlocutors' mental states.³ When it comes to understanding communication, what is important is each interlocutor's *perspective* on the

³ It is often suggested that features of the context of utterance play a role in *determining* the content of what is said, or more broadly communicated. For discussion and criticism of this claim, see Bach (2012).

context of utterance (where these perspectives are just one facet of the context of utterance itself). This is a cognitive category, comprising information – knowledge, beliefs, experiences, etc. – that interlocutors bring to bear in their production and interpretation of an utterance. In principle, any information could be relevant here, but common factors include beliefs about the earlier conversation leading up to the exchange, about an interlocutor's habits, preferences, goals, and motivations, about an interlocutor's social identity, about social institutions, about the physical environment in which the utterance is produced, about stereotypes and 'lexical effects' associated with expressions used in the utterance, and about conversational principles or norms (see Bezuidenhout, 2002, p. 177; Cappelen and Lepore, 2005). Call this category 'contextual information' or 'contextual assumptions'. Where the interpretative resources, and knowledge of the language, required to correctly interpret an utterance are largely shared across interlocutors (either prior to the utterance, or as a result of information introduced by the utterance), we can think of this as 'common ground' in Stalnaker's (1978, 2002) sense. Stalnaker does not require that interlocutors' perspectives on the common ground align perfectly (2002, p. 701). However, in this paper, I am primarily concerned with communication in which there are more significant discrepancies between these perspectives – that is, cases in which there is little common ground between participants. Stalnaker (1978) treats these as 'defective contexts'.⁴ As we will see, defective contexts are especially common in certain kinds of online environment.

Before turning to online communication, however, let us first consider the role of contextual information in utterance interpretation more generally. An agent's utterance of a sentence on a particular occasion can be interpreted in indefinitely many ways. Bach (2005, p. 16), for example, claims that it is platitudinous to say that what an agent means in uttering a sentence is underdetermined by the linguistic meaning of that sentence. Some authors have even suggested that it is more or less impossible for agents to be fully explicit when expressing themselves in language – leaving some room for interpretation is inevitable (Carston, 2009; Bezuidenhout, 2002). The audience must thus rely on their perception of the context of utterance to interpret an assertion. This reliance is needed to resolve a range of interpretative issues that are left open by the surface form of the utterance itself – for example, to disambiguate expressions, to assign

⁴ For a discussion of the prevalence of 'defective contexts' in relation to the common ground framework, see Peet (2021).

Epistemic Bubbles and Contextual Discordance

referents to proper names and indexicals, to recover the explicature of the assertion, to recover any implicatures intended by the agent, and to recognise which further inferences the agent intends her assertion to license. Interlocutors must also rely on contextual assumptions when choosing what to utter: they must select a signal that, in combination with available contextual information, will give their audience a good chance of recovering their intended message, or at least something close to it. Here is an example:

1) There is milk in the fridge.

When an agent asserts (1), there is an enormous number of more or less similar contents that she might be trying to communicate.⁵ It is not simply that her audience must work out which fridge is being referred to. In one context, the agent might be trying to communicate that there is enough fresh milk in the fridge in the kitchen in which the utterance is produced to make one medium sized latte. But, in another context, she might be trying to communicate that there is still spilled milk on the fridge shelf that has yet to be cleaned. Interlocutors do not usually specify with such precision what they intend to communicate in the surface form of the utterance. Rather, they rely on their audience to use their knowledge of the context of utterance to fill out these details in their interpretation.

Although it may be hidden by its banality, a quite astounding amount of shared contextual information is needed for communication to succeed in the simple example just given.⁶ Interpretation of any utterance requires a general understanding of relevant conversational norms and some kind of commonsense person theory (Fricker, 1994), which agents acquire and refine over a lifetime. But, on top of this, a great deal of domain-specific knowledge is involved. Communicating information about milk that could be used in a latte requires that interlocutors possess shared (or relevantly similar) beliefs about the ways milk can be used, about types of coffee-based beverage, about how much a human typically drinks in one sitting, about why certain items are refrigerated, and also about the interlocutors' particular habits and preferences (especially those relating to coffee-making). Consider, for example, that in a

⁵ This example is also discussed in Pollock (2023) and Carston (2009).

⁶ Plausibly, communicate success is a graded, or perhaps even context-sensitive, notion (see Bezuidenhout, 1997; Pollock, 2020). However, we can talk loosely of communication succeeding simpliciter when the degree of success is high.

social milieu in which it was unheard of to use milk in coffee, an utterance of (1), in a coffee-making context, could seem like a non-sequitur, and it would be difficult to work out an agent's intended message and its implications without asking her what she meant.

Despite the sheer volume of information required for successful interpretation, it is often fairly easy for an audience to work out roughly what their interlocutor intended to communicate in these everyday examples. Even in such familiar contexts, however, things can go awry. Perhaps the agent didn't realise that her audience was seeking plant milk, or perhaps she didn't realise that she was making coffee for 6 people, or that she has different standards for when milk can be considered fresh. Even when the interlocutors are quite aware of one another's differing standards, this is no guarantee that utterance interpretation will run smoothly. You know that your housemate is aware of your standards for milk-freshness – you have argued about this many times – but is she accommodating you by adopting your standards, or is she protesting by using her own? This is to say, it is not always clear what information is, or ought to be, considered relevant to interpretation. When communicating in-person, about everyday issues, we can often discover and correct misunderstandings as they become apparent – 'Sorry, I forgot you were vegan!', or 'I meant fresh by *my* standards – yours are unreasonably strict!'.⁷ Many of the cases of online communication that social epistemologists are interested in, in contrast, are not like this simple example in a number of crucial respects: online environments more often involve communication between interlocutors with mismatched contextual assumptions, and these differences can be perpetuated by reduced opportunities to provide missing contextual information or to repair defective exchanges. In the following subsections, I describe these obstacles in more detail.

3.1 Contextual Discordance

In our online communities, it is possible for interlocutors to share a great deal of common ground. This may be the case when we are communicating online with our friends and family, or with members of interest or hobby groups. Other online environments, however, present us with serious obstacles to communication. Agents are often attempting to communicate (or unwittingly communicating)

⁷ See Drożdżowicz (2022) for discussion of how misunderstandings can be repaired.

Epistemic Bubbles and Contextual Discordance

with others in situations in which relatively few contextual assumptions are shared with their audience. This may happen, for example, when agents share public posts on social media: a public tweet can in principle be viewed by each and every one of Twitter's 368 million monthly users (and, in fact, also by internet users without a Twitter account).⁸ Often users of social media will be communicating with people who lead very different lives, and have very different beliefs and experiences, from them. This, in itself, need not undermine communication: successful communication does not require that interlocutors share most of their beliefs or agree with one another about norms for interpretation. Rather, it requires only that interlocutors are aware of one another's perspectives (those that are relevant to interpretation of the utterance). To return to the example from earlier: I don't need to agree with your standards for milk-freshness when interpreting your tweet of (1); I do, however, need to recognise which standards are relevant to interpretation of the tweet. The problem is that, in practice, audiences often know relatively little about the perspectives of those whose posts they consume. This may happen when the producer of an utterance is a stranger to her audience, or chooses to remain anonymous, or simply because audiences consume so many tweets that it would be infeasible for them to acquire information about the authors, even when this information is readily available. Social media platforms like Twitter further compound these issues by maintaining a restrictive limit of 280 characters per tweet, which makes it difficult to provide much information in the utterance itself. Tweets are sometimes parts of longer threads, which may provide more contextual information; however, only a very motivated audience would take the time to read an entire thread.

On the other side of the communicative equation, too, the producer of a public post often has such a large potential audience that she could not, in practice, acquire knowledge about the interpretative dispositions of each member of her audience, nor could she easily select a sentence that is well-suited to communicating her message to all of them at once even if she did have this information. Moreover, it is sometimes the case that an agent does not set out to be understood by every interlocutor who might view her post (and with good reason): as Frost-Arnold (2021) points out, members of marginalised groups and counterpublics may create online spaces for communication, organising, and community-building, which employ local norms that are not intended to be taken up by those outside of the

⁸ This figure is as of 2022 (Dixon, Nov 15, 2023).

community (see also, Habgood-Coote, [forthcoming](#); Ashton, 2020). A lack of common ground between two users of a social media platform, then, is not necessarily a bad thing, all-things-considered; my point, rather, has been that it is a *common* occurrence in online contexts which, for better or worse, will undermine communication between many interlocutors.

In addition to a mere lack of common ground, our online contexts can be especially prone to ‘collapse’ in ways that undermine communicative success (Record and Miller, 2022; Frost-Arnold, 2021). Context collapse is a concept employed in sociology and media studies. Roughly speaking, it is the phenomenon whereby different contexts blur or merge into one another (Meyrowitz, 1985; boyd, 2008; Davis and Jurgenson, 2014). The term ‘context’ here is used in a different way to the way it is used when referring to a context of utterance. Davis and Jurgenson (2014, p. 477) present their notion of context as follows, ‘we define context in terms of role identities and their related networks. Encompassing space, place, history and situation, context refers to the identity meanings activated through interaction with a particular social network.’ Contexts, on this understanding, are organised in relation to different identities or roles of the agent. For example, an agent may have roles as both a rock climber and a philosopher and, in performing each role, is related to different (yet possibly overlapping) networks of agents, who will have different (yet possibly overlapping) expectations of how they will behave. Working with this understanding of context, Davis and Jurgenson define context collapse as ‘the overlapping of role identities through the intermingling of distinct networks’ (2014, p. 477) They distinguish between two varieties: context ‘collisions’, in which two or more contexts accidentally collapse into one another, and context ‘collusions’, in which the collapse is intentional.

The notion of a context of utterance differs from the role-focused notion just introduced in that it takes an utterance, rather than an identity, as its locus. Nonetheless, I think we can usefully apply the notion of context collapse to contexts of utterance in ways that illuminate our understanding of online communication. Without glossing over too many subtleties, we can think of the role identities of interlocutors as one significant aspect of the context of utterance (broadly construed) – as already noted, the context of utterance includes interlocutors’ beliefs and expectations of one another, and these attitudes are, in turn, affected by role identities. Note that, where there is ‘context collapse’, in the sense just introduced, it does not follow that the context of utterance itself is thereby collapsed. Rather, I suggest we think of context collapse as contributing

Epistemic Bubbles and Contextual Discordance

potentially conflicting or mismatched information to the context of utterance – for example, conflicting beliefs of the agent and her audience, conflicting norms for interpretation, or simply increased uncertainty regarding these factors. For example, suppose our rock-climbing philosopher tweets the following: ‘Just solved a problem I’ve been working on for weeks!’ This agent may intend this message to be interpreted as referring to a problem with a research paper she is writing, but her audience may not share her perspective. Some members of her audience may (reasonably) take her to be talking about a bouldering problem, others may be unsure as to how to interpret her. In this case, the agent and her audience may form incompatible or mismatched beliefs about the intended content of the assertion as a result of differing expectations relating to the different role identities of the speaker. These beliefs form part of the context of utterance itself.

The phenomenon of context collapse is not the same as the mere absence of contextual information or common ground. As such, context collapse presents somewhat different obstacles to successful communication. Even in cases in which interlocutors share a great deal of background beliefs, it may be unclear which of these beliefs are relevant to utterance interpretation in cases of context collapse. The issue here is that the same agent will follow different norms for interpretation, or take different information to be relevant to interpretation, in different environments, or amongst different groups. Where there is context collapse, we should expect either that interlocutors will make different assumptions, or that they will be presented with multiple conflicting assumptions in ways that make interpretation more difficult: it will be unclear which assumptions are in play, or which norms ought to be prioritised.⁹ To take an artificial example, suppose that an agent’s social network consists entirely of bankers who love fishing. In this situation, the agent and her interlocutors might have a great deal in common in terms of background beliefs, expectations of one another, perceptions of their physical environment, and even (we can stipulate) upbringing and life experiences. Still an audience member in this community will have trouble interpreting an utterance of ‘We’re meeting by the bank at noon’, without further information.

Davis and Jurgenson (2014) note that all contexts involve some degree of collapse. However, the phenomenon is especially prevalent

⁹ It may be useful to use Roberts’ (2012) notion of a ‘Question Under Discussion’ to frame this point: where there is context collapse, it may be unclear which QUDs are in play, or what order they appear in the stack.

in certain kinds of online environment. This is in large part due to the fact that the volume of role identities that can collide around the production and interpretation of a single online utterance is often far greater than in an offline conversation.¹⁰ In what follows, I will refer to both a lack of shared context, and the consequences of context collapse (in the role-identities sense), as ‘contextual discordance’. Contextual discordance can undermine communication in a wide variety of ways. Some of the problems posed by context collapse in online communication are discussed in detail by Record and Miller (2022) (see also Frost-Arnold, 2021).¹¹ They point out that context collapse can make it harder to recover implicatures (Record and Miller, 2022, p. 7); and they connect this issue to the literature on norms of assertion, arguing that context collapse can lead to differences or instability across interlocutors’ expectations or beliefs about the norms for assertion and for post-sharing. This, in turn, can affect attributions of credibility and responsibility. For example, as Rini (2017) argues, there are not yet established norms for when posts (and retweets in particular) constitute endorsements and, as a result, agents will often attempt to deny responsibility for problematic content that they share.¹² As such, contextual discordance can obscure what kind of speech act an agent is performing. Record and Miller’s focus here is the online spread of what they call ‘epistemically toxic content’ rather than epistemic bubbles. However, the points they make about the role of context collapse in undermining communication are relevant to understanding the challenges posed by online communication more broadly. I would add that context collapse (and other forms of contextual discordance) can have problematic consequences beyond the perception of norms and the recovery of implicatures: contextual discordance can make it harder to determine even the literal or conventional meaning of the agent’s assertion by making it unclear how to resolve disambiguation, reference assignment, and the like. It can make it harder to grasp the explicature or intended meaning of an utterance. And it can make it harder to draw correct or intended

¹⁰ This is not a necessary difference between online and offline communication. A book, for example, has an enormous potential audience. However, books are quite unlike social media posts in many of the further respects identified in this section.

¹¹ Frost-Arnold (2021) also discusses further epistemic consequences of context collapse.

¹² For further discussion of this issue see Arielli (2018) and Marsili (2021).

Epistemic Bubbles and Contextual Discordance

inferences beyond implicature. These various interpretative obstacles can combine to create quite a mess.

3.2 Opportunities for Repair

As I have just explained, communication in certain kinds of online environment tends to involve a higher degree of contextual discordance. A second factor that contributes to difficulties communicating online is that there are often fewer opportunities for providing missing or clarifying contextual information. Consider that, in contrast with in-person communication, posts on social media platforms often lack indications of tone or force, which some agents would otherwise rely on to interpret an utterance.¹³ For example, because many posts are at least partly text-based, they will often lack non-verbal or non-lexical cues that could aid interpretation, such as body language and intonation. This means there are often fewer contextual cues packaged into the communicative event itself that might provide the scaffolding needed to build an interpretation that bridges gaps or differences in contextual information. Moreover, even when such cues *are* present – for example, in a TikTok video – their efficacy can also be undermined by context collapse: body language and other non-verbal cues are not universal and so, just as with text-based posts, it can be difficult for an agent to correctly identify which norms for interpretation are in play.

In addition to this lack of ‘on-board’ contextual information, there is also often comparatively little opportunity for further dialogue between interlocutors in which misunderstandings can be corrected. This is not to say there is no opportunity for dialogue, of course – one can reply to a post to ask for clarification. However, it may be less common for an audience to seek such clarification in relation to social media posts and it would not be feasible for the author of a post to provide such clarification to every member of her audience who might want or need it. It is worth emphasising here that, even in in-person communication, the process of repairing a context is often challenging. Peet (2021) identifies a range of factors that can prevent the repair of defective contexts. For example, interlocutors may refuse to accommodate presuppositions, or fail to challenge perceived misunderstandings, and they may not even recognise that a context is in need of repair in the first place – many

¹³ This is not to say that only online communication lacks such indicators.

misunderstandings go unnoticed (cf. Drożdżowicz, 2022). Peet argues that it is often difficult to identify and repair defective contexts in communication in general and, as such, these defective contexts might be the norm for communication. For the reasons explained above, we should expect this issue to be especially pronounced in online contexts where opportunities for dialogue are reduced, or simply not utilised.

Online contexts are not uniformly contextually impoverished in comparison with offline contexts: there are even sources of contextual information that are typically only available online. For example, it is common for Twitter users to list several identity categories that they belong to in their ‘bios’; this information can be very useful in interpreting a user’s tweets – it provides evidence as to which of the user’s role identities may be especially relevant to interpretation – and would not usually be available when communicating in person. In addition, a user’s entire post history is often public, whereas no such records of offline utterances are commonly available. Despite this availability, however, for the reasons explained above, it is unlikely that this information will be accessed by the majority of users who view a tweet: there simply is not enough time to gather contextual information for all of the posts one consumes, even when this information is easily accessible. The sheer volume of utterances (and interlocutors) that we are confronted with online thus plays multiple roles in generating and sustaining contextual discordance: it is responsible, both for contributing an enormous amount of conflicting contextual information to the context of utterance via context collapse, and for simultaneously preventing agents from resolving resulting contextual discordance due to the practical constraints imposed by the limits of our time and attention.

In this section, I have presented a range of ways in which contextual discordance can pose obstacles to communication, and explained how this problem can be particularly pronounced and hard to repair in certain kinds of online environment. In the next section, I argue that the problems posed by contextual discordance suggest that attempting to add missing information or sources to an epistemic bubble will do very little to shatter it.

4. Online Communication and Epistemic Bubbles

An epistemic bubble, recall, is a social epistemic structure in which relevant information is simply missing. Nguyen suggested that shattering a bubble can be relatively easy: one simply adds the missing

Epistemic Bubbles and Contextual Discordance

information or source and thereby improves coverage reliability. He writes,

... epistemic bubbles are relatively fragile. Relevant sources have simply been left out; they have not been discredited. It is possible to pop an epistemic bubble by exposing a member to relevant information or arguments that they have missed. (Nguyen, 2020, p. 145)

However, I think that merely being exposed to information is often insufficient to enable one to escape an epistemic bubble in any meaningful sense. The problem here is not the mistrust or discrediting of sources that is characteristic of echo chambers. Rather, the problem is that an epistemic bubble shapes the knowledge that is possessed by the agent within it and thus constrains her interpretative resources. When attempting to interpret a new utterance, the agent can appeal to no more than the existing knowledge that she has at her disposal, and the information given in the communicative exchange itself. In what follows I will argue that, as a result, even when a new source of information is added to an epistemic bubble, this is no guarantee that the information they provide can be integrated into the epistemic perspective of the audience who is exposed to it – this information often remains inaccessible to her.

Epistemic bubbles, as we have seen, tend to be populated by people with whom we have much in common. Indeed, it is this feature of them that is thought to be epistemically problematic: epistemic bubbles connect us predominantly with those who share our perspective on the world; they thus tend to reinforce this perspective, rather than challenging or expanding it. Because of this, agents within our bubble tend to be similar to us with respect to their speech and interpretative dispositions – for these dispositions are grounded in our epistemic perspectives. What this means is that it is relatively easy to communicate with others within an epistemic bubble because we tend (or are more likely) to share contextual assumptions with these people: we have enough common ground for communication to run smoothly – at least with respect to certain domains. This relationship is often asymmetric in online communication. That is, your familiarity with the epistemic perspective of someone you follow on Twitter might enable you to interpret her utterances very accurately – the two of you share beliefs about how *she* ought to be interpreted; but this interpretative skill may not be mirrored back at you, especially if this person does not follow you back or take an interest in what you have to say. Bubble-membership itself is plausibly often asymmetric in this way.

Due to the relationship between epistemic bubbles and interpretative resources, then, an agent will typically be relatively good at interpreting utterances from sources within her bubble. Things are quite different, however, when the agent reaches outside of a bubble. When she does so, she will find herself in unfamiliar contextual territory: that is, she will be in contextual discordance with her interlocutor. In these cases, the agent lacks the background knowledge required to accurately interpret what her interlocutor is trying to communicate – either at the level of literal meaning, explicature, implicature, speech act, inference, etc. When this happens, although the agent is ‘exposed’ to relevant missing information – she observes the interlocutor’s utterance – this information is nonetheless still inaccessible to her in some sense (or to some degree). She may recover the wrong content, derive unintended implicatures, or draw incorrect conclusions due to misunderstanding. As a result, she may employ the content she recovers in reasoning in ways that are counterproductive or contrary to the intentions of its source. Importantly, agents who make this kind of mistake are not being irrational. Rather, they may be employing their new beliefs in reasoning as best they can, given the epistemic position they are working from. The problem is that communication is simply very hard in the absence of shared or concordant contextual assumptions and, by the very nature of epistemic bubbles, reaching outside of them requires communicating in the absence of this common ground to a greater or lesser degree.

There are two lessons I would like to draw from this. The first is that contextual discordance presents a special challenge to escaping epistemic bubbles. I return to this issue in Section 5. The second lesson to draw, more briefly, concerns the role of epistemic bubbles in explanations of the so-called ‘post-truth’ phenomenon. Nguyen (2020) claims that epistemic bubbles are not especially useful in explanations of why agents reject claims for which there is overwhelming evidence. He writes,

Notice that epistemic bubbles alone cannot explain the post-truth phenomenon. Since epistemic bubbles work only via coverage gaps, they offer little in the way of explanation for why an individual would reject clear evidence when they actually do encounter it. Coverage gaps cannot explain how somebody could, say, continue to deny the existence of climate change when actually confronted with the overwhelming evidence. One would be tempted, then, to accuse climate change deniers of some kind of brute error. But echo chambers offer an

Epistemic Bubbles and Contextual Discordance

explanation of the phenomenon without resorting to attributions of brute irrationality. (Nguyen, 2020, p. 151)

I agree with Nguyen, of course, that epistemic bubbles cannot wholly explain this phenomenon on their own. However, I do think that they may have a somewhat greater role to play than he allows for here. That is, even a good faith attempt to interpret new evidence can fail if one's interpretative resources lead one to misunderstand the content of this evidence, or what its implications are, etc. Linguistic barriers to escaping echo chambers are briefly discussed by Nguyen. Following Jamieson and Cappella (2008), he describes how those who maintain echo chambers use their own private language, in which familiar terms are given alternative meanings and new jargon is introduced in order to strengthen the in-group's identity (Nguyen, 2020, p. 146). My point is that, even in a rather benign situation, in which the agent is not in the grip of any conspiracy theories, or subject to manipulation, she can still fail to properly *integrate* the information that is presented to her if the interpretative resources supplied by her epistemic bubble are insufficient.

5. Escaping Epistemic Bubbles?

The argument in the preceding has been that, whilst it may be relatively easy to expose oneself to missing sources of information, and thus achieve a kind of coverage reliability, it is much harder to properly digest this information in ways that allow it to play a meaningful role in our epistemic lives. What does this mean for our prospects for escaping an epistemic bubble? One obvious direction for the resolution of this problem is to implement strategies for increasing common ground between interlocutors. Record and Miller (2022, p. 14ff.) suggest some solutions to the problem of epistemically toxic content, which employ this kind of strategy. They suggest, for example, encouraging the introduction of stable norms that govern responsibility for the content of social media posts, and for the interpretation of ambiguous elements of online interfaces (such as emojis). They also identify a range of ways to add more context to posts. For example, they suggest that platforms could employ algorithms that position related posts next to each other in a newsfeed, such that contextual information is collated rather than fragmented; they also consider that platforms can make it easier for (or even require) users to add context to their own posts and reposts. Frost-Arnold (2021, pp. 445–6) also considers proposals for redesigning social media platforms (from Hull *et al.*, 2011; and McMillan

Cottom, 2016), which may aid in counteracting context collapse. However, she is sceptical that such technological solutions would be sufficient, suggesting also that individual users need to change their online practices to develop habits of trustworthiness and discretion (Frost-Arnold, 2021, p. 450).

Both Record and Miller, and Frost-Arnold, are addressing somewhat different challenges to the one that I focus on here. Nonetheless, as these challenges all broadly concern contextual discordance, it is natural to wonder whether strategies like these could be employed to improve our prospects for shattering epistemic bubbles. However, while I do think that it is useful to strategize ways to improve communication in online environments (in cases where this improvement would be all-things-considered preferable), I think it is unhelpful to think of epistemic bubbles as structures that can be shattered or burst, even when additional contextual information is added. This way of thinking about epistemic bubbles simply reproduces a naïve model of communication as floating largely free of contextual considerations. To see this, first consider the initial picture of epistemic bubbles with which we began: on this picture, an epistemic bubble is a structure that can easily be shattered if only the agent within it is exposed to information from sources that had previously been omitted. I argued that it is inappropriate to describe the agent's bubble as 'shattered' in such circumstances because agents communicating across the boundaries of a bubble will lack the interpretative resources required to properly understand an utterance from the new source; and, because of this, they will fail to integrate the new information into their epistemic perspective.¹⁴ For these same reasons, however, I don't think it is helpful to think of an epistemic bubble as something that can be burst, if only we offer *a bit more* contextual information to the agent to support integration of the original content – this picture, too, fails to properly acknowledge the role of contextual information in communication. The main issue is that, especially in online environments, agents are often communicating in circumstances in which contextual discordance is very pronounced. As such, the amount of contextual information that a platform would need to provide to render the context non-defective is extremely high, and the further an agent reaches outside of her bubble, the greater this problem is likely to be. A related

¹⁴ As it stands, this claim is potentially consistent with the idea that the agent may nonetheless gain testimonial knowledge in these circumstances. However, for an argument for the claim that this lack of understanding can undermine testimonial knowledge acquisition, see Pollock (2021).

Epistemic Bubbles and Contextual Discordance

consideration is that utterance understanding is something that comes in degrees. In practice, there is no limit to how well, or how much better, we could understand an assertion from another interlocutor. There is always more contextual information that could be given to an agent to enrich her understanding in ways that bring to light new aspects of her interlocutor's perspective on the world. Moreover, there is very often a deeper or subtler understanding of a proposition possible than that which the agent arrives at: a deep understanding of a proposition from a domain of technical or practical expertise can be a lifetime of work for an individual. Most all agents develop an incredibly rich understanding of propositions in particular domains, of course. But this deep understanding is not the kind of thing that an agent could hope to achieve with respect to all, or even many, of the domains of information that are important or relevant to her life. There is simply too much information in the world, and not enough time to truly comprehend all of it.

With the preceding in mind, I want to propose an alternative model of epistemic bubbles than the one presented by Nguyen. Nguyen's account sees an epistemic bubble as comprising a network of sources, and the information they share, which an agent is connected to via the observation of utterances. On this approach, information can exist within your bubble regardless of how well you understand it; thus, bubbles are relatively easy to burst. Whilst it may be useful to include this network in a model of an agent's broader epistemic environment, I think that an account of epistemic bubbles should reflect the limitations that our interpretative resources place on our access to, and comprehension of, the information we consume. In order to demarcate what information is available within a bubble, I propose appealing to the notion of common ground. On this second approach, although the members of your epistemic bubble may be just those with whom you are connected via utterances, the *information* available is limited to that which is expressed with utterances that you are able to correctly interpret.¹⁵ The idea is that, as two agents may share common ground with respect to one domain, but not another, the information within your bubble is not just any information shared by the members of your network, but only the information shared via utterances that your common ground with these members enables you to understand. Given the nature of utterance understanding, noted above, this is plausibly a matter of degree:

¹⁵ Thank you to Patrick Connolly for suggesting ways to clarify this issue.

information exists within your bubble to the degree that you are able to correctly interpret the utterances used to express it.

If we adopt this alternative model of epistemic bubbles, the result is that attempting to escape one's bubble is not merely a matter of adding missing sources, exposing oneself to new information, and thereby improving coverage reliability. For example, the fact that you follow climate scientists on Twitter does not guarantee that the information they share regarding their research is available within your epistemic bubble. If you lack the interpretative resources to understand their utterances, then following these agents on social media, and even reading their posts regarding climate science, is just one small step in the direction of genuinely integrating the information they share into your own epistemic perspective. This point, of course, does not merely apply to science communication on social media. Similarly, adding activists or political pundits to your social networks does not automatically entitle you to consider all of the information they share as falling within your bubble – the information available to you will depend on which utterances your common ground enables you to interpret, and this will vary across different domains of discourse. To capture new information within your bubble, depending on your starting point, may require significant work, building common ground that can provide the scaffolding for improved utterance interpretation. And this is plausibly work that often involves gradual progress over time. In contrast to the 'fragile' conception of epistemic bubbles that authors like Nguyen have been working with, then, I suggest we instead adopt an 'elastic' conception: that is, epistemic bubbles cannot be shattered, but are instead slowly and incrementally expanded to improve our (inevitably partial) perspective on the world, and our ability to understand others within it. This is not a project that we could hope to complete simply by improving the composition of our social networks, or adding a bit more context to social media posts (although these things surely help); rather, it is a potentially unbounded project that, if we work very hard, we can make modest progress towards within a lifetime.

6. Conclusion

In this paper, I have argued for a particular view of the relationship between epistemic bubbles, on the one hand, and interpretative resources, on the other: epistemic bubbles mark the boundaries of our knowledge of the world and, as such, we can appeal to no more

Epistemic Bubbles and Contextual Discordance

than is available within a bubble in our attempts to understand utterances concerning domains beyond it. For this reason, attempts to ‘shatter’ an epistemic bubble by adding missing sources of information will often be ineffective in practice: communicating across the boundaries of a bubble is communication in the absence of common ground and, without this common ground, the likelihood of misunderstanding is very high. This problem will be especially pronounced in online networks, where contextual discordance between interlocutors is rife, and where there are fewer opportunities to repair defective exchanges. It may be possible to improve online communication by providing more contextual information to aid in the interpretation of utterances from unfamiliar domains. However, I have suggested that it is unhelpful to think of epistemic bubbles as structures that we can, in practice, burst or escape simply by adding more information or sources thereof. Rather, adding contextual scaffolding to online contexts will help us to, gradually and incrementally, enhance our epistemic perspective on the world, and our ability to understand others, but this perspective will always be partial. Thought of this way, the problem with the structure of our online networks is not that they keep us trapped within our epistemic bubbles – for that is inevitable – but that they keep us from expanding our epistemic horizons.

Acknowledgements

This work was supported by the Research Council of Norway (FRIPRO, project no. 315066). I am very grateful to Fintan Mallory, Austin Baker, Drew Johnson, Hugo Mota, Ole Sjøflot, Harish Pedaprolu, Mirela Fuš, Anna Drożdżowicz, Alex Davies, and Patrick Connolly for comments in conversation and on earlier drafts of this paper.

References

- Emanuele Arielli, ‘Sharing as Speech Act’, *Versus*, 127:2 (2018), 243–58.
- Natalie Ashton, ‘Why Twitter is (Epistemically) Better Than Facebook’, *Logically* (November 26, 2020), accessed April 2024 at <https://www.logically.ai/articles/why-twitter-is-epistemically-better-than-facebook>.
- Kent Bach, ‘Context ex Machina’, in Zoltan Szabo (ed.), *Semantics versus Pragmatics* (Oxford: Oxford University Press, 2005), 15–44.

- Kent Bach, 'Context Dependence (such as it is)', in Manuel Garcia-Carpintero and Max Kölbel (eds.), *The Continuum Companion to the Philosophy of Language* (London: Continuum, 2012), 153–84.
- Anne Bezuidenhout, 'Pragmatics, Semantic Underdetermination and the Referential/Attributive Distinction', *Mind*, 106:423 (1997), 375–409.
- Anne Bezuidenhout, 'Truth-Conditional Pragmatics', *Philosophical Perspectives*, 16 (2002), 105–34.
- danah boyd, *Taken Out of Context: American Teen Sociality in Networked Publics*, Ph.D. thesis, University of California, Berkeley, CA (2008).
- Robyn Carston, 'The Explicit/Implicit Distinction in Pragmatics and the Limits of Explicit Communication', *International Review of Pragmatics*, 1:1 (2009), 35–62.
- Jenny L. Davis and Nathan Jurgenson, 'Context Collapse: Theorizing Context Collusions and Collisions', *Information, Communication & Society*, 17:4 (2014), 476–85.
- Stacy Dixon, 'Number of X (formerly Twitter) Users Worldwide from 2019 to 2024', *Statista* (November 15, 2023), accessed April 2024 at <https://www.statista.com/statistics/303681/twitter-users-worldwide/>.
- Anna Drożdżowicz, 'Making it Precise—Imprecision and Underdetermination in Linguistic Communication', *Synthese*, 200 article no. 219 (2022).
- Elizabeth Fricker, 'Against Gullibility', in Bimal Matilal and Arindam Chakrabarti (eds.), *Knowing from Words* (Boston: Kluwer, 1994), 125–61.
- Karen Frost-Arnold, 'The Epistemic Dangers of Context Collapse Online', in Jennifer Lackey (ed.), *Applied Epistemology* (Oxford: Oxford University Press, 2021), 437–56.
- Sanford Goldberg, 'If That Were True I Would Have Heard About It by Now', in Alvin I. Goldman and Dennis Whitcomb (eds.), *Social Epistemology: Essential Readings* (Oxford: Oxford University Press, 2011), 92–108.
- Joshua Habgood-Coote, 'Towards a Critical Social Epistemology of Social Media', in Jennifer Lackey and Aidan McGlynn (eds.), *Oxford Handbook of Social Epistemology* (Oxford: Oxford University Press, forthcoming).
- Gordon Hull, Heather R. Lipford, and Celine Latulipe, 'Contextual Gaps: Privacy Issues on Facebook', *Ethics and Information Technology*, 13:4 (2011), 289–302.
- Kathleen Hall Jamieson and Joseph Cappella, *Echo Chamber: Rush Limbaugh and the Conservative Media Establishment* (Oxford: Oxford University Press, 2008).
- Neri Marsili, 'Retweeting: Its Linguistic and Epistemic Value', *Synthese*, 198:11 (2021), 10457–83.
- Tressie McMillan Cottom, 'Twitter's New @Replies Re-design Isn't Just Stupid; It's Really Stupid', *Medium* (October 30, 2016), accessed April 2024 at <https://tressiemcphd.medium.com/twitters-new-replies-re-design-isn-t-just-stupid-it-s-really-stupid-c471ca254f0a>.
- Joshua Meyrowitz, *No Sense of Place: The Impact of Electronic Media on Social Behavior* (New York: Oxford University Press, 1985).

Epistemic Bubbles and Contextual Discordance

- Boaz Miller and Isaac Record, 'Justified Belief in a Digital Age: On the Epistemic Implications of Secret Internet Technologies', *Episteme*, 10:2 (2013), 117–34.
- C. Thi Nguyen, 'Echo Chambers and Epistemic Bubbles', *Episteme*, 17:2 (2020), 141–61.
- Eli Pariser, *The Filter Bubble: What the Internet is Hiding from You* (London: Penguin, 2011).
- Andrew Peet, 'Defective Contexts', in Justin Khoo and Rachel Sterken (eds.), *Routledge Handbook of Social and Political Philosophy of Language* (New York: Routledge, 2021), 193–208.
- Joey Pollock, 'Context and Communicative Success', in Tadeusz Ciecierski and Pawel Grabarczyk (eds.), *The Architecture of Context and Context-Sensitivity* (Cham: Springer, 2020), 245–63.
- Joey Pollock, 'Linguistic Understanding and Testimonial Warrant', *Erkenntnis*, 88:2 (2021), 457–77.
- Joey Pollock, 'Testimonial Knowledge and Content Preservation', *Philosophical Studies*, 180:10 (2023), 3073–97.
- Isaac Record and Boaz Miller, 'People, Posts, and Platforms: Reducing the Spread of Online Toxicity by Contextualizing Content and Setting Norms', *Asian Journal of Philosophy*, 1:41 (2022).
- Regina Rini, 'Fake News and Partisan Epistemology', *Kennedy Institute of Ethics Journal*, 27:2 (2017), E43–E64.
- Craige Roberts, 'Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics', *Semantics & Pragmatics*, 5:6 (2012).
- Robert Stalnaker, 'Assertion', in Peter Cole (ed.), *Syntax and Semantics 9: Pragmatics* (New York: New York Academic Press, 1978), 315–32.
- Robert Stalnaker, 'Common Ground', *Linguistics and Philosophy*, 25:5 (2002), 701–21.
- Cass Sunstein, *#Republic: Divided Democracy in the Age of Social Media* (Princeton: Princeton University Press, 2018).

JOEY POLLOCK (joanna.pollock@ifikk.uio.no) is a senior lecturer in philosophy at the University of Oslo. Her research focuses on the nature of communication, and the intersection between issues in philosophy of language and social epistemology.