

Context Collapse Online

Lucy McDonald

For *Conversations Online*, edited by P. Connolly, S. Goldberg, and J. Saul, forthcoming

PLEASE CITE PUBLISHED VERSION

Introduction

Social media platforms have significantly changed how we interact with one another, in part because they facilitate what internet theorists call ‘context collapse’. This is defined by Alice Marwick and danah boyd as the flattening of multiple, diverse audiences into one single audience (2011, 122), which they think has radical psychological and psychosocial effects. In an ordinary offline, face-to-face conversation, the thought goes, we have a good idea of the audience’s identity and the context we occupy, and therefore of which social norms apply. As such, we can maintain a high level of control over the impression we make to our audience, and can cultivate an appearance of authenticity. As Erving Goffman observes, we vary our self-presentation in different conversations and with different audiences, but this need not lead to tensions or perceived inauthenticity, because these different contexts can be kept separate (1969).

The same is not true of our interactions on social media. There, ‘social contexts we used to imagine as separate co-exist’ (Marwick and boyd 2011, 130) and we are therefore ‘forced to grapple simultaneously with otherwise unrelated social contexts that are rooted in different norms and seemingly demand different social responses’ (boyd 2014, 31). On a given social media platform, one’s audience may consist of school friends, work colleagues, family members, and ex-partners, all of whom might potentially see one’s content. The collapse of these ordinarily separate contexts makes it difficult ‘for people to engage in the complex negotiations needed to vary identity presentation, manage impressions, and save face’ (Marwick and boyd 2011, 123). Offline, we can present different versions of ourselves in different contexts, but social media ‘problematizes the individual’s ability to shift between these selves and come off as authentic or fake’ (Marwick and boyd 2011, 124).

The literature on context collapse gives us a good sense of the phenomenon’s effect on our relationships and our identities. Yet context collapse has ramifications far beyond this. In this chapter, I’ll explore the impact of context collapse on *communication*. Context collapse is surely one reason why online communication seems so different from offline communication, and why it has given rise to new communicative problems. As the internet becomes a central site for civic discourse, it is important that we evaluate the applicability of our standard

pragmatic theories and models to this new communicative domain.¹ In our analyses of offline communication, we assume that our audiences are both known and static, as are conversational norms and goals. On social media, context collapse renders these assumptions inappropriate, and as such we cannot assume that the kinematics of online communication resemble the kinematics of offline communication.

The chapter will proceed as follows. In §1, I give the history of the notion of context collapse. In §2, I draw on David Lewis's scorekeeping framework to propose that we think of communicative context collapse online not as merely the flattening of multiple audiences into one, but rather as the existence of multiple overlapping and intermingling language games, governed by different rules and with different scoreboards. In §3, I consider some of the political, epistemic, and moral ramifications of communicative context collapse online.

1. Social Context Collapse

1.a A Potted History

While context collapse is widely discussed in the context of social media, the concept was first discussed in the context of television, radio, and the telephone by communications scholar Joshua Meyrowitz. In *No Sense of Place* (1986), Meyrowitz draws on the 'situationist' theories of Erving Goffman to explain how electronic media transformed social interactions.

One of Goffman's key insights was that human behaviours are strongly shaped by context; we behave differently depending on where we are and whom we are with. Depending on this context, we accentuate, or 'foreground' certain characteristics, and minimise, or 'background' other characteristics (Goffman 1969). Different contexts come with different social scripts, with different roles for different participants. Goffman develops this idea using a dramaturgical metaphor; each context is like a play, where the stage must be set, actors must utter the right lines and wear the right costumes, technicians must operate the right sounds and lights, and no one who is supposed to stay backstage should stray on stage (1969, chap. 3).

Implicit in Goffman's work, Meyrowitz noticed, was an assumption that contexts are physically bounded. For Goffman, a behavioural region is 'any place that is bounded to some degree by barriers to perception' (1969, 109). Yet electronic media radically undermined the 'traditional relationship between physical setting and social situation' (Meyrowitz 1986, 7). For example, two people talking on the telephone are participating in a shared context, but each occupies a different physical setting. During the call, certain scripts apply, and the interlocutors take up distinctive roles, but these features of the context are not indexed to or prescribed by the physical setting.

¹ The Pew Research Centre has described social media as 'the new public square', writing that it is 'increasingly home to civil society, the place where knowledge sharing, public discussions, debates, and disputes are carried out' (Smith et al. 2014).

Goffman had also assumed, Meyrowitz notes, that there is a ‘relatively stable social order with fixed rules, roles, social occasions, and institutions’ (1986, 2). Goffman seemed to think this followed from the physical boundedness of contexts; if physical settings determine the social order, and physical settings are not easily nor quickly changed, then the social order will not be easily or quickly changeable either.² But electronic media has changed this, too:

Perhaps one of the reasons that theorists of situations and roles have tended to view social situations as relatively stable is that it is extremely rare for there to be a sudden widespread change in walls, doors, the layout of a city, or in other architectural and geographical structures. But the change in situations and behaviors that occurs when doors are opened or closed and when walls are constructed or removed is paralleled in our times by the flick of a microphone switch, the turning on of a television set, or the answering of a telephone. (Meyrowitz 1986, 39–40)

Electronic media, Meyrowitz argues, has produced ‘a very discernible rearrangement of the social stages on which we play our roles and a resulting change in our sense of “appropriate behaviour”’ (1986, 4). It did this by destroying physical boundaries on contexts, and thereby leading to the intermingling of previously bounded social situations, creating ‘potentially infinite degrees and patterns of situation overlap’ (1986, 42). It produced entirely new situations with new scripts and roles, and forced participants in these situations to adapt.³

Meyrowitz illustrates this dimension of electronic media with the example of Stokeley Carmichael. Carmichael was a civil rights activist, heavily involved in the Black Power movement, who frequently engaged with both Black and white audiences. Given many white people were very hostile to the Black Power movement, Carmichael communicated differently with the two groups. Yet when his speeches began to be broadcast on TV and radio, he found himself addressing both audiences simultaneously; the two contexts, as it were, collapsed. He could no longer carefully tailor his rhetoric to different audiences to avoid alienating either. In the end, he addressed himself primarily to Black people. This decision, Meyrowitz observes, ‘filled his secondary audience with hatred and fear and brought on the wrath of the white power structure’ (1986, 43).

Internet scholar dana boyd, who is widely credited with popularising the concept of context collapse, drew heavy inspiration from Meyrowitz but turned her focus away from TV, radio, and the telephone and onto social media platforms like Facebook and Twitter (boyd 2002, 2010, 2013, 2014; Marwick and boyd 2011). These are online communities in which users can create and share various forms of content with one another. On these platforms, users create personal profiles consisting of identifying information, like their name and location, and curated content, like photos and text posts. They can connect with (by ‘friending’, ‘following’ or ‘subscribing to’) other personal profiles, thereby developing complex online networks.

² The idea here is not that the social order cannot be changed – rather, it is that social reorganisation requires, among many other things, radical architectural and geographic changes.

³ Meyrowitz does not think context collapse is unique to electronic media; he also thinks events like weddings facilitate context collapse, as they typically feature a ‘combination of many different audiences’ (1986, 5). He stresses, though, that electronic media produce a more extreme kind of context collapse. Weddings are typically one-off, and one can leave them, but in modern society one cannot really escape electronic media (Ibid).

Connecting with a person in this way gives a user access (to varying degrees, depending on the platform and their privacy settings) to that person's profile and content, and creates various public and private ways of interacting with them.

boyd argues that context collapse on social media is produced by the architecture, or more specifically the 'affordances', of these platforms (boyd 2010; Marwick and boyd 2011). One distinctive affordance of social media, boyd argues, is *persistence*; online content is 'automatically recorded and archived', which facilitates asynchronous conversation and makes online exchanges much less ephemeral than, for example, a face-to-face spoken interaction (boyd 2010, 46–47). Another is *replicability*; online content is very easily reproduced, much more so than any other form of communication. Because content is also easily modifiable online, 'content can be transformed in ways that make it hard to tell which is the source and which is the alteration', such that 'what is replicated may be altered in ways that people do not easily realize' (2010, 47).

A third affordance of social media is *scalability*; like television and other electronic media, but perhaps to an even greater degree, social media enables users to distribute content incredibly widely, such that there is always a possibility of 'tremendous visibility' (2010, 48). Finally, social media platforms also facilitate *searchability*; accessing information and finding other users on them is quick and easy.

boyd writes that these affordances 'do not dictate participants' behaviour, but they do configure the environment in a way that shapes participants' engagement' (2010, 39). One effect of them, she argues, is context collapse; social media forces us to 'contend with groups of people who reflect different social contexts and have different expectations as to what's appropriate' (2010, 50). This occurs because the persistence, replicability, scalability, and searchability of social media platforms create situations in which we are unsure of our audience's identity; 'with the audience invisible and the material persistent, it is often difficult to get a sense of what the context is or should be' (2010, 51). Even when we know who our immediate audience is (or who they are likely to be), we know that there is a potential audience 'far greater and from different contexts' (2010, 50).

In a paper with Alice Marwick, boyd argues that context collapse makes it difficult to construct an authentic online persona (Marwick and boyd 2011). Authenticity, boyd and Marwick think, requires a stable, singular identity but on social media we find ourselves grappling with conflicting social scripts. We might be expected to interact in different ways with friends, colleagues, and family, for example, but all of these groups are potential audience members online, and we cannot satisfy all norms at once. If one tries to alternate between adherence to these conflicting norms, one will appear chaotic at best and disingenuous at worst.⁴

⁴ Jenny Davis and Nathan Jurgenson distinguish between two kinds of context collapse: context collusions and context collisions (2014). Context collusion is deliberate; it is 'the process whereby social actors intentionally collapse, blur, and flatten contexts' (2014, 480). This might be done, for example, during disaster relief efforts to distribute information and connect volunteers. Context collision, meanwhile, is what happens when 'different social environments unintentionally and unexpectedly come crashing into each other' (2014, 480–1).

It is interesting to consider whether context collapse on social media, as analysed by Marwick and boyd, is different from context collapse produced by TV and radio, as analysed by Meyrowitz. I wager that it is, in terms of both scale and kind. Firstly, most of us typically consume TV and radio shows, but we do not produce them. This means that we do not have to grapple with the task of tailoring our material to suit a variety of audiences; this challenge is faced only by the broadcasters and content-producers themselves. On social media, however, we are all broadcasters and content-producers, and it is therefore very difficult to escape the range of difficulties generated by context collapse. boyd makes a similar point, observing that while the dynamics of social media are not new, ‘they were never so generally experienced’ (2010, 49).

Secondly, social media users also experience a more personal kind of context collapse. Marwick and boyd point out that whilst TV and radio broadcasters typically address ‘a faceless mass’, social media users address an audience which is ‘unidentified but contains familiar faces; it is both potentially public and personal’ (2011, 129). Thus on social media there is a ‘presumption of personal authenticity and connection’ (Ibid). This might raise the stakes for alienating members of one’s audience; one risks damaging or even losing extremely significant personal relationships. Since social media content is also much less mediated than TV and radio content, in so far as it is not edited or checked by a production team, and you need no substantial economic, cultural, or social capital to produce it, mistakes and missteps are even more likely on social media than they are on radio or TV.

1.b New Directions

Marwick and boyd’s work has illuminated the psychological and psychosocial ramifications of context collapse online, and in particular the challenges it poses for how we present ourselves to others and manage our relationships. In this chapter, however, I will focus on a different dimension of context collapse. I will look not at its social ramifications, but rather at its *communicative* ramifications. Context collapse creates challenges not just for our relationships, but also for how we communicate. It challenges us to think carefully about what we say and how to say it in a way which does not alienate our diverse audience, and increases the risks of our words being misunderstood or taken out of context, of interruptions, diversions, and abuse, of shaming, and of passionate disagreements when interlocutors from very different communities cross paths.

Several philosophers have already drawn attention to the communicative ramifications of context collapse. Though he does not use the term ‘context collapse’ itself, Sanford Goldberg argues that the indeterminate contours of audiences in online conversations makes it hard to delineate conversational context (2021). He observes that a user often cannot tell who is part of their ‘conversation’, as some users may be silent, and it is not always clear when users have entered or left a conversation. Moreover, it’s not obvious that the contours of conversational participants online are determinate in the first place: it is hard to say whether a person who makes a quick comment in a thread and then moves on, never returning to it again, is a current

participant in that conversation. In addition, online conversations can be extended over long periods of time – discussions can be revived by a user many years after they were initiated. For all of these reasons, Goldberg thinks, it can be difficult to discern the context set of an online conversation – we cannot tell which propositions are mutually assumed by conversational participants.

Both Karen Frost-Arnold (2021, 2023) and Isaac Record and Boaz Miller (2022) argue that by making sufficient contextual information harder to obtain, context collapse facilitates misunderstanding online. Frost-Arnold begins with the claim that to settle the meaning of propositions online, users must draw on what Elizabeth Fricker calls a ‘knowledge context’ (Fricker 2012, 66), which is a representation of the conversational context, including information about the interlocutors and the goals and norms of the conversation (Frost-Arnold 2021). Frost-Arnold argues that when a post is seen by users who are not its target audience, a possibility made very likely by context collapse, those users are unlikely to share a knowledge context with the original poster, and as such are unlikely ‘to recover the proposition that the speaker intended to convey to their imagined audience’ (2021, 444). This can lead those users to form false beliefs about what the poster said. In the same vein, Record and Miller argue that that to understand content on social media platforms, users must fill in ‘information gaps’ concerning the post’s meaning, but this process is impeded when multiple interpretations of the post are offered by multiple audiences, making it hard to hit upon the correct interpretation (2022).⁵

This work demonstrates some of the challenges context collapse poses to interlocutors in online spaces. It also demonstrates the limited applicability of standard pragmatic models to online discourse. Work in pragmatics often assumes that in a given conversation, all interlocutors are aware of one another’s presence and identities, that the person who hears an utterance is also the person to whom it is addressed, and that interlocutors are generally cooperative, contributing to an evolving common ground and working towards a common goal.⁶ These idealisations often don’t apply to offline conversation, but online they rarely ever apply, in part due to context collapse.⁷ As such, we need a new way of modelling online discourse. My goal in the rest of this paper will be to develop such a model, and use it to deepen our understanding of the ways context collapse shapes conversations online.

2. Communicative Context Collapse

⁵ In addition to this philosophical work, linguists have also looked at context collapse and communication. Jannis Androutsopoulos explores how context collapse affects language style (2014), and Philip Seargeant, Caroline Tagg, and Ngampramuan explore how context collapse affects language choice and code-switching among bilingual Facebook users (2012).

⁶ On the idealisations and assumptions of contemporary philosophy of language, see Beaver and Stanley (2018) and Dever and Cappelen (2019).

⁷ For example, Ellefson (this volume) shows how the presence of trolls in online spaces radically undermines the assumption that participants in a conversation have shared communicative goals. Trolls, she thinks, exploit the shared conversational goals of others in order to pursue their own private goals, disrupting the conversation in the process.

2.a Communicative Contexts

In this section I will canvas some different ways of conceiving of communicative context collapse on social media, using tools from philosophy of language. First, I must clarify how I will be understanding the notion of ‘context’. There is a general tendency in the literature on context collapse to equate contexts with audiences, yet it seems social media actually creates context collapse in a more expansive sense: it brings together not only different audiences, but also different sets of norms, concepts, conversational styles, priorities, and stores of knowledge.

According to one popular way of thinking in pragmatics, a conversation is best thought of not as a mere assembly of interlocutors, but as a language game with a set of rules and a scoreboard. Wittgenstein first introduced the notion of a *language game* to philosophy (1953), proposing that language games are rule-governed activities whose rules are collectively agreed upon and then upheld by its players. These rules determine what kinds of move one can make, as well as the meaning of any given word or act: ‘the meaning of a word is its use in the language’ (1953, §43). In this sense, communicating is like playing a game. And there are different kinds of these language games, with different norms and different purposes, which contribute to different ‘forms of life’ (1953, §19, §23).

The kind of language game one is participating in will determine which conversational contributions one can make. We can understand the modality of this ‘can’ in at least two ways, drawing on John Searle’s distinction between constitutive and regulative rules. The *constitutive* rules of a language game determine which conversational moves are possible. These rules ‘constitute (and also regulate) forms of activity whose existence is logically dependent on the rules’ (Searle 1964, 55). If you are playing football, there is a constitutive rule according to which kicking the ball into the other team’s net (provided some other conditions are met) counts as a goal; the act of goal-scoring is made possible by this constitutive rule of football. Similarly, if you are a judge in a court room, there is a constitutive rule according to which your declaring the defendant guilty makes it the case that they count as guilty in law. This move is made possible by the rules of the language game, and cannot be performed in other language games, like a casual conversation with a friend. The *regulative* rules of a language game, meanwhile, do not determine which moves are possible, but rather stipulate how moves *ought* to be performed; they ‘regulate activities whose existence is independent of the rules’ (Ibid). Violating these rules in a language game will make one’s utterance hard to understand or vulnerable to sanction on account of being thought irrelevant, impolite, immoral, or illegal, but the utterance will still be recognised as a move within that conversation nonetheless.

For some kinds of conversations there are thought to be conversational norms (constitutive or regulative) that apply by default and endure throughout such conversations. These might be institutional in nature; for example, the rules of the House of Commons dictate that members cannot address other members as ‘you’, but rather must address them through the Chair.⁸ Or these norms might be cultural; in a conversation between participants of a shared cultural

⁸ *Rules of behaviour and courtesies in the House of Commons*, House of Commons, September 2021, <https://www.parliament.uk/globalassets/documents/rules-of-behaviour.pdf>

background, there may be certain topics one cannot talk about (or conversely that one must talk about), or there may be established patterns for turn-taking.

H.P. Grice also shows that there are some general conversational norms that appear to be in play in most conversations. Grice's Co-operative Principle holds that conversational participants should aim to make their contributions 'such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which [they] are engaged' (1975, 45), and will be expected to do so. This amounts to requirements that contributions be informative, truthful, relevant, and clear.

Other times, what one can and cannot say will be determined not by enduring, pre-established conversational norms but rather by the evolving 'scoreboard' of that conversation. We owe this concept of a conversational scoreboard to David Lewis, who extends the language game analogy by proposing that for each particular conversation there is a scoreboard, like one might find in a game of baseball (1979). This scoreboard keeps track of all of the moves made by each interlocutor in the game, including all speech acts performed and attempted, as well as other phenomena, like permissibility facts, referents, and facts about salience.

This scoreboard determines which moves in the conversation might count as felicitous, or to use Lewis's language, 'correct play', but it also determines many other phenomena, including the referents of the definite descriptions and proper names one uses, one's point of reference when one uses directional language and indexicals (what it means to 'come here' will depend on where 'here' is), and the truth of predications of vague predicates (which depends on standards of precision already established in the conversation). It also contains what Craige Roberts calls 'the question under discussion', which is the focus to which conversational participants are currently oriented, against which conversational moves will be interpreted (2015).

Here is an example of how the scoreboard of the conversation to which one is contributing shapes and constrains one's utterances. Imagine that you and I are discussing US politics, and in particular the policies of Joe Biden. We repeatedly use 'the President' to refer to Biden. Imagine now that I wish to talk instead about South Korean politics. It would be inappropriate for me to start using 'the President' to refer to Yoon Seok-youl without making clear that I am doing so, because we have established a temporary conversational norm that Biden is the most salient president in this context, and that 'the President' refers to him. Hence previous moves in the conversation determine which future moves (in this example, moves involving definite descriptions) are permissible.

This isn't to say that we cannot change such rules. Indeed, if it becomes plain to you that I am now talking about South Korean politics, I may succeed in changing the salience ranking of the conversation so that 'the President' does now refer to Yoon Seok-youl. This highlights an important disanalogy between language games and sports games: moves in language games can be subject to processes of what Lewis calls *accommodation*. Whereas in a sports game there are some moves which will never count as correct play, a scoreboard in a conversation can automatically update to make an infelicitous move ultimately felicitous (1979).

Thinking of online contexts as language games with complex sets of rules and evolving scoreboards helps us identify three possible ways of understanding ‘context collapse’ in online communication: context collapse as participating in an unknown language game, context collapse as participating in a language game with unknown and changing participants, and context collapse as attempting to participate in a language game while other, overlapping language games occur simultaneously around you. I will present and evaluate each of these in turn, concluding that the latter is the best way of thinking about communicative context collapse online.

2.b Communicative Context Collapse Online

Perhaps context collapse online makes it such that we cannot know which language game we are participating in. Hence perhaps it is like turning up on a sports field, and jumping into a ball game half-way through, knowing neither the rules of the game nor the score so far.

We can flesh out this initial analogy in two different ways. It could be that the player is aware of the rules of different possible games, but does not know which of those games she is currently participating in. For example, she could know how to play both baseball and cricket, but be unsure which of these she is playing. This captures some of what it feels like to communicate online. If you jump into a long Twitter thread halfway down, not only will you not know what the general rules of the conversation are, but without access to the previous discussion, you will also have no idea what the conversational score is. You won’t know what the QUD is, whom the names used refer to, what the salient standards of accuracy are, et cetera. This makes missteps highly likely.

This first construal of the analogy entails that once the player figures out which game she is playing, as well as the score, she will be able to participate reasonably smoothly and avoid violating conversational norms. Yet this is not quite what online communication feels like; our lack of knowledge often seems more severe than this. A second, perhaps more accurate form of the analogy holds that when the player turns up on the sports field they lack not the knowledge of which of a set of recognisable games they are playing, but rather knowledge of the games themselves. Often when we are communicating online, it is not that we cannot distinguish which of a set of conversational games we are playing, but rather that the game itself is unfamiliar to us.

Think back to what it feels like to first join a social media platform; not only do you not understand the interface and the moves it makes possible, but you will also not yet be au fait with the norms governing such moves. For example, you may not know what clicking the heart symbol next to a tweet does, nor when or why one should do so. For someone new to internet communication entirely, the communicative devices used – from abbreviations (*btw*, *nbd*, *imho*, *nsfw*, *afaik*) to neologisms (*retweets* and *upvotes*) to reaction GIFs – can be inaccessible.

Most of us have now been using such platforms for years, so the platform-specific language may no longer pose a challenge. But we will still often find ourselves in conversations where

we are unaware of the norms; we may ask questions that we are not supposed to ask in such a conversation, or use words which have a conversation-specific meaning we are unaware of. This will indeed feel similar to joining a ball game at the park that one does not know the rules of.

And yet this analogy still does not seem quite right, because it assumes that the player knows who her fellow players are. On the field, you can see who you're playing with, because of their proximity, the arrangement in which they stand, and their behaviours. Indeed, in face-to-face conversation you have similar information; as Sanford Goldberg notes, there are a range of conventions 'that enable us to discern the contours of the conversational participants' (2021, 179). Online, you cannot know who your fellow language game players are, in several ways.

Online, there may be some potential players who are invisible to you, because they are not currently speaking. They may be overseeing, moderating, or simply 'lurking' in the conversation, possibly shaping its norms and trajectory in ways that are not clear to you. And even among those people visible to you, you cannot know who is playing the game and who is not. It can be hard therefore to distinguish between players, observers, and uninvolved bystanders. Of those online players who are visible (because they are creating, sharing, and responding to content), you may see merely a username, perhaps a photo. Some may be entirely anonymous. As such their identity is often obscured in a way it might not be in a real life sports game.

Moreover, as discussed earlier, Goldberg points out that the contours of online conversations are less determinate than the contours of offline conversations (2021, 179). Offline, there are clear conventions signalling the beginning and end of the conversation, as well as for individual agents to enter and exit that conversation. Online, it is unclear when conversations begin and end – indeed, some seem to drag on endlessly – and participants drop in and out. As such there is rarely a fixed set of players in the language game.

Let us therefore consider a modified analogy. Perhaps context collapse online involves not only not knowing the rules of the language game you are playing, but also not knowing who you are playing with – in the senses that of the people visible to you, you don't know who is a player; that you can never be sure of the existence of invisible players; and that of those players visible to you you may sometimes know very little about them. To modify the analogy, online communication is like joining an enormous game on the sports field, where some players are wearing a mask and non-players are strolling amidst the players, in a way that makes it hard to identify the players. Some players may be facing away as if they are not playing, but may jump in and out of the game unexpectedly.

So how does this affect communication? I noted before that participants in a conversation bring their own norms and expectations to the conversation, and what may be appropriate to say to person A may not be appropriate to say to person B. Hence one could easily make cultural and social missteps if one were to say the wrong thing to the wrong person. More generally, not knowing exactly who is listening can make it hard to know what to say and how to say it, in a way that does not alienate at least some of the audience. Knowledge of one's fellow players is

also required to play the game appropriately; for example, if the audience is large, intimate disclosures may be inappropriate. If the audience is small, they may be permissible.

And yet this modified analogy is still not quite right, for at least two reasons. Firstly, it assumes that the given online communicator is participating in just one language game, and that the same is true of everybody else. As Goldberg notes, the contours of online conversations are indeterminate. As such, wherever and whenever one is communicating, multiple language games are likely bleeding into one another. It is hard to underestimate the scale and diversity of online conversation.

Secondly, the modified analogy assumes good faith on the part of the other players, i.e., that they recognise that they are norm governed and are attempting to follow those norms. Anyone who has spent any time on the internet knows that there are some people online whose motivation is to wreak havoc and/or cause suffering. Sometimes they do this by deliberately and openly violating conversational norms. Other times they do this by manipulating and bending norms, exploiting others' expectations that they will abide by those norms. These users are often known as 'trolls'. Gretchen Ellefson (this volume) argues that trolls often exploit and undermine the shared goals of other participants in an online conversation in order to further their own unshared, and often concealed, goals. Patrick Connolly (2022), meanwhile, argues that often a troll's utterance may appear to its target to be a sincere contribution to the conversation, but it is in fact performed with a view to creating entertainment for onlookers outside of that conversation.⁹

Thus I propose the following scenario as the best analogy for communicative context collapse online. Imagine you are in the park, playing what you think and hope is a game of baseball. The park is very busy, and other people around you are also playing various ball games, like rounders, cricket, baseball, and softball. Mix ups inevitably ensue; a ball you pitched in your game of baseball may end up near the group playing cricket, and they may count that as an innings in their game. You might lose track of which game you are playing, and make a move which is permissible in cricket, but not in baseball. Some people in the park may deliberately wander from game to game, attempting to participate in each, sometimes even simultaneously, with varying degrees of success. Some may seek to confuse and upset games; they might walk around throwing balls in all directions, for example, or a cricketer might masquerade as a baseball player. Rivalries might arise between the cricketers and the baseballers.

Onlookers in the park may also lose track of what is occurring, and might chastise players for misguided reasons; for example, they might chastise your baseball fielder for wearing gloves because they think she is playing cricket, a game in which fielders are not supposed to wear gloves. The cricket fans may come to idolise you as a champion cricketer, when you identify as a baseball player. You may even start trying to play baseball in a way that appeals to those who think you are playing cricket.

Online communication often involves multiple overlapping language games occurring simultaneously, with indeterminate contours and an everchanging set of participants, only ever

⁹ For more on trolls, see Barney (2016), Cohen (2017), DiFranco (2020), and Morgan (2022).

partially known to the other players. Different contexts collapse into one another not only in the sense that different audiences come together, but also in the sense that there are multiple scoreboards and multiple distinct sets of evolving norms.

3. The Impact of Communicative Context Collapse

Now that we have settled on an apt analogy for context collapse online, we are better placed to assess its effects on online discourse, (political, epistemic, and moral). It is important to stress, though, that ascertaining the exact effects of this phenomenon is in part an empirical matter. Without further empirical testing and verification the below discussion should be seen as speculative rather than conclusive.¹⁰

3.a Political Ramifications

In an optimistic mood, we might think that context collapse brings us closer to the ideal deliberative public sphere outlined by Jürgen Habermas (1992), by exposing deliberators not only to a diverse range of people and views, but also to a diverse range of communicative practices.

Greater exposure to alternative viewpoints may give us more evidence that those who defend alternative ways of life may have legitimate grounds for doing so, and this may make for more knowledge, and more democratised, productive discourse which is more likely to yield decisions which promote the common good. Greater exposure to alternative communicative practices, meanwhile, may offset a phenomenon Iris Marion Young calls ‘internal exclusion’, whereby actors formally included in deliberative discourse nonetheless find themselves informally excluded by norms of political communication (2002, chap. 2).

Young argues that general norms of articulateness, dispassionateness, and orderliness function to exclude those who make claims in a non-linear fashion and communicate emotionally or ‘disruptively’. She argues that deliberation should make room for non-standard communication like greeting, rhetoric, and narrative. Context collapse makes agents subject to multiple overlapping and conflicting communicative norms, thereby disrupting the hegemony of standard norms like articulateness, dispassionateness and orderliness and perhaps making deliberation less exclusionary.

¹⁰ For data on these issues, see ‘Social Media and Political Dysfunction: A Collaborative Review’, a compilation of research crowdsourced by Jonathan Haidt and Chris Bail. https://docs.google.com/document/d/1vVAtMCQnz8WVxtSNQev_e1cGmY9rnY96ceYuAj6C548/edit. The challenge for internet scholars interested in a specific phenomenon, like context collapse, is finding data which examines the effects of solely that phenomenon. Social media interfaces are extremely complex, and it is difficult to isolate one single feature of these platforms and identify the effects of that feature alone.

In a less optimistic mood, however, we might worry that context collapse in fact makes deliberation harder and more exclusionary. After all, democratic deliberation requires a common vocabulary and a common goal – to identify and promote the common good – that context collapse undermines by introducing a variety of conflicting conversational goals.

That said, perhaps this kind of deliberation is undesirable anyway. Nancy Fraser argues that a single Habermasian deliberative public sphere will inevitably come to be dominated by those with social power, and she defends in its place the *proliferation* of publics, where less powerful agents can escape restrictive and oppressive communicative norms and ‘invent and circulate counterdiscourses, which in turn permit them to formulate oppositional interpretations of their identities, interests, and needs’ (1990, 67).¹¹ She offers as an example the feminist counter publics that enabled women to ‘invent new terms for describing social reality’, thereby expanding ‘discursive space’ (Ibid). Jenny Odell similarly points to the important role closed, small meetings played in the organisation of the Montgomery bus boycott; had these meetings been larger and open to the public, they would have been vulnerable to detraction, derailing, and police harassment (2019, 162).

The internet has in many ways facilitated such counter publics, by making it easier to find others with your identities, interests and needs. Yet when these groups and communities spring up in the public areas of social media – for example, in non-private discussions on Twitter or open groups on Facebook – they may struggle to ever really escape the watchful eyes of dominant social groups, and may find themselves held to the very norms they have actively disavowed. Remaining publicly accessible enables these groups to attract like-minded people, but context collapse makes it likely that their discussions will be shared with unintended audiences, and renders them vulnerable to misinterpretation, intrusion, and derailing. If Fraser’s critique of the Habermasian ideal is correct, then context collapse could still be a foe to democratic deliberation.¹²

The hijacking of anti-racist discussions in the Black Lives Matter movement online illustrates how context collapse creates opportunities for dominant social actors to actively interfere with counter publics focused on the emancipation of marginalised groups. Those who assert ‘Black lives matter’ on social media take themselves to be engaged in a discourse where one of the central questions under discussion is ‘Do Black lives matter?’ The widespread mistreatment and murder of Black people across the United States makes plain that for many people, Black lives do not matter. Black Lives Matter activists seek to change this. However, these activists are often met online with the response that ‘All lives matter’. This response, Jessica Keiser argues, presupposes that there is a different question under discussion – not ‘Do Black lives matter?’ but rather ‘Which lives matter?’ (2021). The response thereby attributes an implicature to those who assert ‘Black Lives Matter’ that they did not make – that white lives do not matter. Keiser claims that this ‘stirs up confusion and divisiveness by systematically distorting

¹¹ Alison Jaggar (2004) and Karen Frost-Arnold (2021) argue that closed communities are valuable not just for political reasons but also for epistemic reasons; they facilitate the creation of knowledge.

¹² It is perhaps for this reason that some groups ‘go private’ with tightly monitored and regulated Facebook groups and Discord servers, thereby creating for themselves a safe haven from context collapse.

affirmations of the value of Black lives as expressions of threat to the value of the lives of others' (Keiser 2021, 8478).

We can see from this example that context collapse makes it easier for racists and their ilk to achieve at least two things. The first is to disrupt and derail anti-oppressive discourse, forcing activists to defend and explain themselves. The second is to drum up support for themselves and their racist projects. By misrepresenting the nature of anti-racist discourse, the All Lives Matter folk can dupe onlookers into regarding them as reasonable and virtuous. Imagine that a Twitter user who does not harbour any explicitly racist beliefs, but who is also not au fait with contemporary anti-racist activism, comes across an online conversation where a user has responded to 'Black lives matter!' with 'All lives matter!' If she at this point lacks any knowledge of what the QUD was for the Black Lives Matter activists who initiated the discussion, she may consider the 'All lives matter!' response quite apposite. As a consequence, she may even align herself with the community who frequently makes such responses, thereby undermining anti-racist work. This is exactly what the 'All lives matter' folk want, and context collapse helps them achieve it; it makes it easier for them to co-opt the conversational context of anti-racist discourse to recruit oblivious onlookers to their cause.¹³

In addition, context collapse also seems to incentivise anti-democratic, anti-egalitarian agents to disguise their communication in a way that makes holding them responsible for harmful and wrongful speech more difficult. To show how this happens, I will draw on Jennifer Saul's work on dogwhistles (2018). Saul follows Tali Mendelberg in thinking there is something like a 'norm of racial equality' in contemporary American political discourse (Mendelberg 2001). This norm does not actually proscribe racism simpliciter, but rather requires lip service to 'an extremely thin sort of formal equality' (Saul 2018, 365) and proscribes very overt racism, like the use of slurs or support of obvious discrimination. When such a norm is in play, racists must get creative; they must find ways to propagate racist ideas to both fellow racists and non-racists in a way that is not seen to violate the norm of racial equality. One solution is the use of dogwhistles.

Dogwhistles are a kind of coded language, and they can take different forms. Some have a secret meaning for hearers 'in the know' and an innocuous meaning to others. Saul gives as an example the emphatic criticism by Republican politicians in the early and mid 2000s of the Supreme Court decision in *Dred Scott v. Sandford* (1857). In this case, the Court denied American citizenship to people of Black African descent, a decision later reversed by the passage of the Thirteenth and Fourteenth Amendments in 1865. The Republicans' public condemnation of the 1857 decision 'officially' signals their opposition to slavery and racism. Yet, insidiously and unofficially, it also signals opposition to abortion. It does this by making salient the possibility that Supreme Court decisions can be reversed. By critiquing *Dred Scott*, Republicans could signal that they would like to reverse another well-known Supreme Court decision: *Roe v. Wade* (1973). Regrettably, this is a feat they subsequently achieved.

Other kinds of dogwhistle subconsciously influence the thinking of hearers in pernicious ways. An example of such a dogwhistle is 'inner city', which when heard by individuals with pre-

¹³ I thank the editors of this volume for helpful discussion on this point.

existing racist attitudes raises those attitudes to salience and makes the individuals more inclined to support racist policies. Unlike those who engage in explicit hate speech, users of dogwhistles benefit from plausible deniability (Saul 2018, 361). If accused of deliberately using the terms with the intention of propagating oppressive and harmful ideas, they can claim that they were unaware that these terms were dog-whistles, and that they thought them wholly benign.

Context collapse incentivises the use of dogwhistles because it makes it more likely that prejudiced speakers will be overheard by those who uphold anti-prejudice norms, and thereby encourages these speakers to be sneakier in their expressions of prejudice to one another. If there is a possibility that your racist utterances will be heard as contributions to a liberal discussion, and that you will therefore be accused of violating the norm of racial inequality, you have good reason to disguise those utterances.

Context collapse also incentivises prejudiced speakers to be sneakier in how they influence outsiders; rather than using explicit propaganda, they are led to use coded language like ‘inner city’ which acts more subtly on hearers’ psyches. Hence we have reason to worry that context collapse has made racist speech (and other prejudicial speech) slipperier, and its peddlers more devious, such that they are harder to blame and punish.¹⁴ As contexts proliferate, and the number of contexts an utterance could be seen as a contribution to increases, prejudiced speakers have incentives to get more and more creative in their camouflage strategies. Online interlocutors have developed an array of sophisticated multi-media dogwhistles, including words, phrases, symbols, images, gestures, and GIFs, which enable users to tailor their coded messages to many different audiences in many different scenarios.¹⁵

3.b Epistemic Ramifications

It is interesting to consider the epistemic effects of communicative context collapse online, too. Context collapse certainly makes it more likely that utterances will be interpreted and evaluated in contexts different from those in which they were first uttered. As such, they will often be interpreted as having different locutionary, perlocutionary, and illocutionary dimensions from the ones the utterer intended and believed them to have in the intended context of utterance, and they will contribute to different conversational scores in unexpected ways. One immediate risk, then, is that context collapse will foster misunderstanding and miscommunication; the original ‘meaning’ of messages will be lost when those messages are transplanted into different contexts, in a virtual version of the ‘telephone game’ played by children. Consequently, agents may end up with a variety of false beliefs (see Frost-Arnold, 2021).

¹⁴ There are of course two sides to this coin. I claim here that it is a downside of context collapse that, by incentivising the use of subtle, sneaky kinds of oppressive speech, it makes it harder to hold people responsible for harmful language. Yet one might think that subtly oppressive speech is all things considered less bad (or less wrong) than overtly, explicitly oppressive speech, and so perhaps it is an *upside* of context collapse that it has incentivised the ‘toning down’ of oppressive speech. I thank the editors for highlighting this possibility.

¹⁵ For more on visual dogwhistles, see Drainville and Saul (forthcoming).

I explained above how racists can manipulate contexts to make it the case that utterances of ‘Black lives matter’ appear to generate an undesirable implicature. I turn now to another way in which utterances are transmogrified in an epistemically significant way due to context collapse online: sometimes a proposition that was clearly true when asserted in context A will be interpreted in context B, where the contextual parameters are fixed in a way that renders the proposition false. This can lead to what Jonathan Jenkins Ichikawa calls ‘contextual injustice’, whereby a speaker finds their utterance unfairly assessed against contextual parameters which are disadvantageous to them (2020, 17).

To modify an example from Ichikawa’s paper, consider female students discussing harassment on Twitter, who assert ‘We know that Professor X is guilty of sexually harassing his students’. In this context, they are using modest, but not unusual, epistemic conditions for knowledge – for them it is sufficient for knowledge that they have received testimony from multiple reliable students about the harassment. Imagine now that another Twitter user interprets the students’ assertion as a statement of criminal guilt, and responds, ‘You cannot know that Professor X is guilty because this has not been proven beyond reasonable doubt’. This immediately raises the bar for knowledge in this context, and renders the students’ assertion false; in this new context with very stringent epistemic standards, they cannot truly ‘know’ that Professor X is guilty.

Ichikawa argues that this employment of contextual parameters can be unfair and disadvantageous to the speakers. Not only might speakers who spoke in good faith nonetheless be accused of lying or being misleading (especially worrying given scaremongering about false accusations of sexual misconduct), but also they might be deterred from further consciousness-raising in ways that set back emancipatory goals. In response to the risk that others will raise contextual standards in ways that threaten their assertions about injustice, speakers may decide to abandon making those assertions entirely.

Ichikawa’s discussion of assertions about harassment assumes the truth of contextualism about knowledge, a thesis which is not uncontroversial, but even if we reject such contextualism, we can still accept the existence of a broader phenomenon whereby agents unjustly change contextual parameters. Ichikawa himself suggests, for example, that a speaker who uses a gradable adjective like ‘good’ or ‘tall’ could become subject to contextual injustice if an interlocutor unfairly changed the context-sensitive parameters for that adjective (2020). Online, this kind of contextual injustice occurs frequently, as agents find their utterances evaluated against contextual parameters they did not believe were in place at the time of utterance. Sometimes agents will deliberately raise contextual standards unfairly, but other times utterances will simply end up in contexts quite different from the ones agents believed they were participating in.

Finally, context collapse may interfere with our ability to weigh the significance of different pieces of information. Nicholas Carr observes that social media ‘trivializes everything’ – ‘a presidential candidate’s policy announcement is given equal weight to a snapshot of your niece’s hamster and a video of the latest Kardashian contouring’ (Carr 2020). Most of us find ourselves in more than one communicative context at any given time, and it can be challenging and jarring to hold in our heads multiple scoreboards, norms, and questions under discussion in a way that enables us to keep track of the all-things-considered significance of different

pieces of information. This is the chaos described in comedian Bo Burnham's 'Welcome to the Internet', in which he sings, 'Welcome to the internet/Put your cares aside/Here's a tip for straining pasta/Here's a nine-year-old who died', characterising the internet as offering you 'A little bit of everything/all of the time' (2021). Similarly, Odell writes of skimming one's Twitter feed that 'Many things in there seem important, but the sum total is nonsense, and it produces not understanding but a dull and stupefying dread' (2019, 159).

A person who is told in a single, clearly demarcated offline conversation that *Roe vs Wade* has just been overturned may consider this news monumentally significant both politically and personally. But when they learn such news whilst simultaneously participating in (and thereby distracted by) conversations about animals, sports, and the best restaurants in London, their ability to discern its significance may be hampered. When serious and unserious conversations are entangled, the risk is that the serious conversations are afforded less appreciation than we might think is personally prudential and morally and politically appropriate.

3.c Moral Ramifications

In our chaotic park scenario, a cricketer may find herself being blamed for violating the rules of baseball, due to confusion as to which game she is participating in. Similarly, context collapse seems to make inapt and/or disproportionate blame more likely in online discourse. An in-joke that someone thought was made among friends in a light-hearted language game might be interpreted by others as a serious and offensive contribution to the language game of political debate. What is not a foul in one game may count as a foul in another.

Sometimes you will make an utterance which you believe satisfies all the norms governing context A, but will be accused of violating the norms governing context B – reasonably or unreasonably. It may be that you should have known you were in context B, or at least you should have known that you would be interpreted as participating in context B, in a way that makes you blameworthy on account of recklessness or negligence. Or it may be that you could not have known or predicted that your utterance would be interpreted in context B, such that any blame that does follow is inapt. Context collapse makes it the case that utterances will be frequently 'taken out of context', understood in a different context from the one in which the speaker believed they were uttered. This can alter their meaning in ways that reflect poorly on the speaker, who may then be accused of misleading, lying, or engaging in disrespectful and/or harmful speech. As such, context collapse effectively envelops the online communicative terrain in intersecting and overlapping trip wires.

As a result, online communicators will often find themselves subject to blame, apt or otherwise, because the proliferation of conversational norms creates a proliferation of means of transgression.¹⁶ Sometimes this blame can be excessive; the combination of multiple

¹⁶ In this way context collapse raises an interesting question about the ethics of praise and blame: under what conditions is it appropriate to hold someone to the standards or rules of a practice or game in which she merely *appears* to be participating?

conflicting norms with the huge audiences available online creates ample opportunities for ‘pile-ons’, or shaming events, where thousands of people publicly blame the target, invite others to join in, and share the news of the purported transgression far and wide. C. Thi Nyugen describes how the ‘Retweet’ function facilitates this:

A couple of unlucky retweets, and suddenly your tweet is splayed across tens of thousands of screens of people who don’t share your background context, who can’t interpret your tweet properly. Jokes are read straight; ironic mockeries of racism are read as actual racism. This kind of thing is so familiar that new media scholars even have a name for it: context collapse. And we know the usual results: misunderstandings, firings, shaming, mob pile-ons. To be sure, not every Twitter pile-on is the result of misunderstanding. Some pile-ons are clear-sighted responses to perfectly legible statements. But many of them involve context collapse. (Nguyen 2021)

Such pile-ons can be deeply psychologically harmful to their targets, and may be wrongful even when the initial blame itself is deserved. They may even be counterproductive, if the effect is to psychologically destroy the target, or force them into hiding, rather than to motivate them to improve themselves.

Conclusion

In this paper I have explored and expanded the concept of context collapse, arguing that in online discourse not only are multiple, diverse audiences flattened into one single audience, but also multiple language games collide, bringing together different players, rules, and scoreboards. Developing David Lewis’s language-baseball analogy (1979), I proposed that communicating online is akin to attempting to play one game or more on a busy field in which numerous other people are playing numerous other games, some similar, some very different. Thinking of communicative context collapse online in this way can help us understand many of the politically, epistemically, and morally interesting features of online discourse.

Though I take no stance on whether context collapse is all things considered desirable or undesirable, it evidently has some undesirable effects, and it is interesting to consider what steps we might take to mitigate those effects. One strategy might be to encourage awareness of the diversity of online contexts, in the hopes that we will recognise contextual manipulation and misinterpretation more easily and exercise more interpretative charity.

Another is to sharply segregate online language games. Just as a solution to the chaotic multi-game park scenario would be to create more distance between the different groups and games, and stricter conditions of entry and participation, online communicators might avoid some of the problems of context collapse (at least on some platforms) by siloing their conversations. Yet the risk of this strategy is that it may lead to what Nguyen labels ‘epistemic echo chambers’ (2020) and may obstruct healthy political deliberation. Deciding on a strategy ultimately requires us to reflect not just on how context collapse harms, but also on what an ideal communicative landscape online would look like, a challenge I leave to future work.

References

- Androutsopoulos, Jannis. 2014. 'Languaging When Contexts Collapse: Audience Design in Social Networking'. *Discourse, Context & Media*, Digital language practices in superdiversity, 4–5: 62–73.
- Barney, Rachel. 2016. '[Aristotle], On Trolling'. *Journal of the American Philosophical Association* 2 (2): 193–195.
- Beaver, David, and Jason Stanley. 2018. 'Toward a Non-Ideal Philosophy of Language'. *Graduate Faculty Philosophy Journal* 39 (2): 503–47.
- boyd, danah. 2002. 'Faceted Id/Entity: Managing Representation in a Digital World'. Cambridge, MA: Massachusetts Institute of Technology.
- . 2010. 'Social Network Sites as Networked Publics: Affordances, Dynamics, and Implications'. In *Networked Self: Identity, Community, and Culture on Social Network Sites*, edited by Zizi Papacharissi, 39–58. New York: Routledge.
- . 2013. 'How "Context Collapse" Was Coined: My Recollection'. <https://www.zephorio.org/thoughts/archives/2013/12/08/coining-context-collapse.html>.
- . 2014. *It's Complicated: The Social Lives of Networked Teens*. New Haven: Yale University Press.
- Burnham, Bo. 2021. *Bo Burnham: Inside*. Netflix.
- Carr, Nicholas. 2020. 'From Context Collapse to Content Collapse'. ROUGH TYPE (blog). 13 January 2020. <https://www.rough.type.com/?p=8724>.
- Cohen, Daniel. 2017. 'The Virtuous Troll: Argumentative Virtues in the Age of (Technologically Enhanced) Argumentative Pluralism'. *Philosophy and Technology* 30 (2): 179–189.
- Connolly, Patrick. 2022. 'Trolling as speech act'. *Journal of Social Philosophy* 53 (3): 404–420.
- Davis, Jenny L, and Nathan Jurgenson. 2014. 'Context Collapse: Theorizing Context Collusions and Collisions'. *Information, Communication & Society* 17 (4): 476–85.
- Dever, Josh and Herman Cappelen. 2019. *Bad Language*. Oxford: Oxford University Press.
- DiFranco, Ralph. 2020. 'I Wrote This Paper For the Lulz: the Ethics of Internet Trolling'. *Ethical Theory and Moral Practice* 23 (5): 931–945.
- Drainville, Raymond, and Jennifer Saul. Forthcoming. 'Visual and Linguistic Dogwhistles'. In *Oxford Handbook of Applied Philosophy of Language*, edited by Luvell Anderson and Ernie Lepore. Oxford: Oxford University Press.
- Ellison, Nicole B., Charles Steinfield, and Cliff Lampe. 2007. 'The Benefits of Facebook "Friends:" Social Capital and College Students' Use of Online Social Network Sites'. *Journal of Computer-Mediated Communication* 12 (4): 1143–68.
- Ellison, Nicole, Charles Steinfield, and Cliff Lampe. 2011. 'Connection Strategies: Social Capital Implications of Facebook-Enabled Communication Practices'. *New Media & Society* 13: 873–92.

- Fraser, Nancy. 1990. 'Rethinking the Public Sphere: A Contribution to the Critique of Actually Existing Democracy'. *Social Text*, no. 25/26: 56–80.
- Fricker, Elizabeth. 2012. 'Stating and Insinuating'. *Proceedings of the Aristotelian Society Supplementary Volume* 86 (1): 61–94.
- Frost-Arnold, Karen. 2021. 'The Epistemic Dangers of Context Collapse Online'. In *Applied Epistemology*, edited by Jennifer Lackey. Oxford: Oxford University Press, 437–456.
- . 2023. *Who Should We Be Online? A Social Epistemology for the Internet*. New York: Oxford University Press
- Goffman, Erving. 1969. *The Presentation of Self in Everyday Life*. London: Allen Lane.
- Goldberg, Sanford C. 2021. 'The Promise and Pitfalls of Online "Conversations"'. *Royal Institute of Philosophy Supplement* 89: 177–93.
- Grice, H. P. 1975. 'Logic and Conversation'. In *Syntax and Semantics 3: Speech Acts*, edited by Peter Cole and Jerry L Morgan, 41–58. New York: New York Academic Press.
- Guerrero-Solé, Frederic, Sara Suárez-Gonzalo, Cristòfol Rovira, and Lluís Codina. 2020. 'Social Media, Context Collapse and the Future of Data-Driven Populism'. *Profesional de La Información* 29 (5): 1–12.
- Habermas, Jürgen. 1992. *The Structural Transformation of the Public Sphere: An Inquiry Into a Category of Bourgeois Society*. Translated by Thomas Bürger and Larry Kert. Cambridge: Polity Press.
- Ichikawa, Jonathan Jenkins. 2020. 'Contextual Injustice'. *Kennedy Institute of Ethics Journal* 30 (1): 1–30.
- Jagar, Alison. 2004. 'Globalizing Feminist Ethic.' In *Setting the Moral Compass: Essays by Women Philosophers*, edited by Cheshire Calhoun. New York: Oxford University Press, 233–55.
- Keiser, Jessica. 2021. 'The "All Lives Matter" Response: QUD-Shifting as Epistemic Injustice'. *Synthese*, 199 (3-4):8465-8483
- Lewis, David. 1979. 'Scorekeeping in a Language Game'. *Journal of Philosophical Logic* 8 (3): 339–59.
- Marwick, Alice E., and danah boyd. 2011. 'I Tweet Honestly, I Tweet Passionately: Twitter Users, Context Collapse, and the Imagined Audience'. *New Media & Society* 13 (1): 114–33.
- Mendelberg, Tali. 2001. *The Race Card: Campaign Strategy, Implicit Messages, and the Norm of Equality*. Princeton, NJ: Princeton University Press
- Meyrowitz, Joshua. 1986. *No Sense of Place: The Impact of Electronic Media on Social Behavior*. New York: Oxford University Press.
- Morgan, Andrew. 2022. 'When Doublespeak Goes Viral: A Speech Act Analysis of Internet Trolling'. *Erkenntnis*.
- Nguyen, C. Thi. 2020. 'Echo Chambers and Epistemic Bubbles'. *Episteme* 17 (2): 141–61.
- . 2021. 'Twitter, the Intimacy Machine'. *The Raven Magazine*. 8 December 2021. <https://ravenmagazine.org/magazine/twitter-the-intimacy-machine/>.
- Norlock, Kathryn. 2017. 'Online Shaming'. *Social Philosophy Today* 33 :187–197.
- Odell, Jenny. 2019. *How to Do Nothing: Resisting the Attention Economy*. New York: Melville House.

- Record, Isaac & Boaz Miller. 2022. 'People, posts, and platforms: reducing the spread of online toxicity by contextualizing content and setting norms'. *Asian Journal of Philosophy* 1 (2):1–19.
- Roberts, Craige. 2015. 'Accommodation in a Language Game'. In *A Companion to David Lewis*, edited by Barry Loewer and Jonathan Schaffer, 345–66. Oxford: John Wiley & Sons, Ltd.
- Saul, Jennifer. 2018. 'Dogwhistles, Political Manipulation, and Philosophy of Language'. In *New Work on Speech Acts*, edited by Daniel Fogal, Daniel W. Harris, and Matt Moss, 360–83. Oxford: Oxford University Press.
- Sergeant, Philip, Caroline Tagg, and Wipapan Ngampramuan. 2012. 'Language Choice and Addressivity Strategies in Thai-English Social Network Interactions'. *Journal of Sociolinguistics* 16 (4): 510-531
- Searle, John R. 1964. 'How to Derive "Ought" From "Is"'. *The Philosophical Review* 73 (1): 43–58.
- Smith, Marc a, Lee Rainie, Ben Shneiderman, and Itai Himelboim. 2014. 'Mapping Twitter Topic Networks: From Polarized Crowds to Community Clusters'. *Pew Research Center: Internet, Science & Tech* (blog). 20 February 2014. <https://www.pewresearch.org/internet/2014/02/20/mapping-twitter-topic-networks-from-polarized-crowds-to-community-clusters/>.
- Vitak, Jessica. 2012. 'The Impact of Context Collapse and Privacy on Social Network Site Disclosures'. *Journal of Broadcasting & Electronic Media* 56 (4): 451–70.
- Wittgenstein, Ludwig. 1953. *Philosophical Investigations*. Translated by G. E. M. Anscombe. Oxford: Blackwell.
- Young, Iris Marion. 2002. *Inclusion and Democracy*. Oxford: Oxford University Press.