Penultimate draft. Please cite published version (when available).

Forthcoming in Conversations Online. Eds. Patrick Connolly, Sanford Goldberg, and Jennifer Saul, OUP.

Imagined Audiences and Common Ground[1]

Karen S. Lewis
Barnard College, Columbia University
klewis@barnard.edu

1 Introduction

Existing theories of conversational contexts are modeled on private, face-to-face conversation as the paradigmatic case. Conversations online differ from face-to-face conversation in significant ways, and it is not clear whether existing theories of conversational dynamics are helpful in making sense of what is going on in communicative interactions online. By 'online conversation' (or, sometimes for variety, 'conversations online'), I mean conversations that take place more or less publicly on social media, in the comments sections of newspapers and blogs, and on discussion forums like reddit. I am not referring to private conversations (such as direct messages or email) that take place online. I am ultimately interested in a full theory of context and conversational dynamics for online conversations, but in this paper, I explore one part of this larger question: (some) features of the common ground.

Following Robert Stalnaker (1978,1999,2002,2014), many popular and fruitful theories of conversational context hold that it includes at least a common ground, if not holding, as Stalnaker does, that the conversational context just is the common ground. The common ground is the shared background of information against which interpretation takes place. It is also dynamic, evolving over the course of a conversation; the very common ground against which interpretation takes place is updated in response to conversational moves, and the updated common ground in turn acts as the background for interpretation for the next conversational move. In the following section I present problems for both the background role and the dynamic role of common ground when it comes to conversations online, based on the observations that conversations online do not involve determinate conversational participants and speech acts are not manifest events, both of which are important to existing theories of common ground. Nevertheless, I argue that there are good reasons for thinking common ground does in fact play both these roles in conversations online, and the way in which it plays these roles is one important element for understanding both what goes right and what goes wrong in conversations online. In section 3, I argue that a speaker in a conversation online has a common ground with an *imagined audience*, rather than an actual audience, and that this imagined common ground plays an important role in successful communication. In section 4, I argue that the imagined common ground does evolve dynamically over the course of the

[1] I'd like to thank participants at the 1st annual Philosophy of Language Association conference in June 2022, a workshop associated with this volume in August 2022, PhLiP in November 2022, the Eastern APA in January 2023, and the words workshop in February 2023 for their comments, questions, and discussion of this paper. Thank you to my research assistants, Katherine Flanagan and Isabel Herburger, for finding many of the examples in this paper and for fruitful discussion.

conversation, but the picture that emerges is not the straightforward one of a single evolving common ground, but rather a series of overlapping imagined and actual common grounds.

2 Problems for Common Ground

Online conversations present *prima facie* challenges to the centrality of the common ground in theorizing about conversational context. I am putting aside, for the purposes of this paper, the question of what is going on in conversations with bots, though it is certainly an important aspect of what is going on in conversations online. For now, I am focusing on what happens when people converse with each other in public conversations on social media using text, images, and recorded videos. While there are platform-specific differences between forums that do matter for how conversations proceed (including issues related to the details of the common ground), the features that present the general problem for common ground are common to all of them. Online conversations raise problems for both of the central roles played by common ground: its role as the background against which interpretation takes place, and its dynamic role as the thing updated by conversational moves. I'll take each problem in turn.

2.1 Problems for the background role

The common ground is the shared background against which conversations take place. A shared background of information raises the question *shared by whom* or *whose information?* Stalnaker famously takes the common ground to be reducible to the propositional attitudes of the conversational participants: "It is common ground that ϕ in a group if all members *accept* (for the purpose of the conversation) that ϕ, and all *believe* that all accept that ϕ, and all *believe* that all *believe* that all accept that ϕ, etc." (Stalnaker 2002, p.716) To define a common ground, therefore, there needs to be a group of people who have the relevant propositional attitudes about the members of that group. Furthermore, requiring such a group is not some peculiarity of Stalnaker's definition. Regardless of what definition one accepts, for a proposition to be common ground is for it to be common ground *for a particular group of people*. In typical face-to-face conversations, in letter and email writing, in texting and instant messaging and the like, the conversational participants are typically well-defined, and so the group of people in terms of which to define the common ground is clear. But when it comes to online conversations, the audience is ill-defined.[2]

Like more traditional public speakers and writers such as journalists, authors, politicians, and celebrities, speakers on social media (i.e. the authors of posts) are faced with *invisible*

[2] Sanford Goldberg (2021) points out a related worry about the determinateness of conversational participants (and the consequent problem of determining the common ground). He argues that unlike offline conversations, there are no clear conventions for signaling initiation of a conversation with a particular audience, or for a participant to signal that they are still part of an ongoing conversation, or that a conversation had concluded. Furthermore, there is both the epistemic problem of knowing who is part of a conversation online (since one is not normally aware of those who read posts but do not comment or react) and the metaphysical problem of determining who is part of the conversation in the first place (e.g. someone who quickly glances at a post and a comment or two but doesn't contribute).

audiences (boyd 2014). An invisible audience is one that is, at least in part, unknown. For any given post on social media, the speaker does not know exactly who (if anyone) will read it. Given the public, highly visible, and spreadable nature of online speech, even the ordinary person with a small social media presence does not know the entirety of the *potential* audience who will see their posts. Further to this, the persistence of online speech – the fact that individual utterances remain more or less permanently available on social media – means that even speakers with relatively few followers or friends at the time of posting might find their posts read quite widely at a later date. So it is unclear how to define a common ground in the first place, since there is not a well-defined audience with whom the speaker shares mutual presuppositions or shared background more generally.

The problem is further exacerbated by the phenomenon of *collapsed contexts*, which arises from having different audiences at the same time, where each audience follows their own social norms. Everyone presents themselves in different ways in different social situations and to different audiences. For example, many people present themselves differently to their colleagues than they do to their childhood friends. A collapsed context occurs when, for example, one runs into a childhood friend at a work event and is required to choose a manner in which to speak and the sort of things to say. The advent of broadcast media increased the occasions for collapsed contexts among public speakers like politicians and celebrities. For example, Joshua Meyrowitz (1985) (cited in boyd 2014) describes the case of the American civil rights activist Stokely Carmichael. When Carmichael began speaking on radio and television, he faced the problem of addressing at the same time both the white political leaders and the southern Black congregations to whom he'd been previously speaking as separate audiences, presenting different talks in different rhetorical styles. In the end, he chose to speak in the rhetorical style he used for Black congregations, which resulted in alienating the white political leaders.

Social media puts all users in the position of the public speaker on broadcast media. The public nature of online speech means that virtually every social media user is faced with collapsed contexts (Marwick and boyd 2011, boyd 2014); posters on social media are speaking at once to different audiences with different social norms. Even those who have relatively few followers or friends face collapsed contexts, as most people have a mix of social groups who read their social media content. Something most of us are familiar with is that social media requires us to navigate speech in a situation in which our potential audience includes (for example) co-workers, distant and close family members, people from high school, friends from our childhood camp, current friends (from potentially different social groups), etc. Furthermore, the spreadable and persistent nature of online speech means that even if a poster is fairly certain about who the initial audience will be, they do not know the extent of the potential future audience.

When media scholars talk about collapsed contexts and invisible audiences, they are concerned with navigating the norms of different social contexts and the presentation of self. (To name but a few of many possible considerations, someone posting on social media has to decide whether to post about serious or fluffy or quirky topics, whether to post personal or

professional photos, whether to emphasize their liberal or conservative political viewpoints, whether to be sarcastic, or emotional, or mysterious, whether to post about their personal life or their hobbies or observations about politics. All such considerations contribute to their online persona on a particular platform at a particular time.) I am not primarily concerned with the presentation of self. I am primarily concerned with conversational contexts in the technical sense – those things philosophers of language and linguists have pinpointed as crucial for interpretation and for tracking information as it flows through a discourse, such as the common ground. Collapsed contexts don't just collapse social contexts, but also potential common grounds for a conversation. Different audiences bring with them different beliefs and different knowledge, and one can take for granted different propositions when addressing, for example, one's childhood friends and one's colleagues. Speakers on social media also have to navigate, what, if anything, they can take for granted as background to a post, given the constant potential for collapsed contexts.

2.2 Problems for the dynamic role

The propositions that are common ground are always changing in response to what happens in a conversation. Changes to the common ground are generally a result of conversational participants reasoning based on the occurrence of manifest events:

> Changes in the common ground, like changes in common knowledge and belief, will normally take place in response to what I will call a "manifest event". A manifest event is something that happens in the environment of the relevant parties that is obviously evident to all. A goat walks into the room, or all of the lights suddenly go out. In such a case, it immediately becomes common knowledge that the event has happened – that there is a goat in the room, or that the lights have gone out. Speech acts will themselves be manifest events (at least when all goes well): when one produces an utterance, it becomes common ground that it has been produced, and when the semantics of the language is common ground, it will be manifest that an utterance with a certain meaning has been produced. (Stalnaker, *Context*, p.47)

Online environments do not lend themselves to making any particular speech act "obviously evident to all". Even if we temporarily put aside the problem of defining who "all" is, it is not clear that a post is a manifest event, in the sense of being *obviously* evident to anybody (except the speaker). Perhaps on certain platforms, people with a large enough following can be fairly certain that the algorithm is such that their post will be seen by many people. Or maybe the existence of notifications on some platforms makes it likely that many people will see such a post. Even in these cases, it is not clear that it counts as a manifest event – is doing something that one can be fairly certain that many people will see the same as doing something that "obviously evident to all"? In any case, for many who post online, there is nothing that guarantees that the fact that a post was made is obviously evident to the relevant parties (if we

can answer the question of who the relevant parties are), or to anyone at all. Posting online is much more like making an assertion by shouting into a loud, crowded auditorium, not knowing if anyone will hear, or who it will be if someone does, than the manifest events described in the above quotation from Stalnaker.

I take each of the problems presented in this section in turn. In the following section, I address the background role for common ground, arguing that there is an important role for an *imagined* common ground both from the speaker's perspective (in message design) and the hearer's perspective (in interpretation). If this is right, then it establishes a role for a *static* notion of common ground in conversations online. In section 4, I turn to the dynamics of common ground, arguing that we can indeed define a common ground that evolves over the course of a single conversation and that it plays a role similar to that of common ground in conversations offline, though complications persist.

3 Common ground with an imagined audience

3.1 Imagined Audiences

As already argued, there is no actual audience at the time of making a social media post. It is not just that there is nobody *there* as is in a face-to-face conversation, but that the audience is invisible and diverse – the speaker does not know who will read their post, nor can they reasonably have just one specific person or group in mind as the audience of a public post[3], unlike some other mediated communication, such as email, texts, or letters. Many who research mediated communication that reaches wide audiences, such as fiction writing, broadcast media, and social media, argue that in these cases speakers write or speak to an *imagined audience,* a speaker's "mental conceptualization of the people with whom [they are] communicating". (Litt 2012, p.1) In section 3.2, I will argue that speakers on social media post against a background of a common ground with an imagined audience, and when all goes well, the actual audience interprets the post against this same common ground. But first I will describe in more detail what I mean by *imagined audience*.

We could think of the imagined audience as operative in all communicative situations, including face-to-face conversations, as there is always a way in which a speaker is conceptualizing their audience. However, in face-to-face conversation, the imagined audience is generally based very closely on who the actual audience is (what the speaker knows about them, how they are behaving, etc.). When it comes to invisible audiences and collapsed contexts, speakers rely on an envisioned audience that is less rooted in the actual audience. As Eden Litt (2012) points out, the degree to which one is free to fictionalize their audience depends on the speaker's social role. The fiction writer is free to fictionalize their audience completely (Ong 1975), whereas newscasters, politicians, and journalists should be informed by various features of their actual past audiences and potential future audiences in imagining their audience (Coddington et al.

[3] Both because there is no reasonable expectation that these would the *only* people who read nor is there in general a reasonable expectation that *every* person in the intended audience would read it.

2021, Nelson 2021). Nevertheless, the imagined audience to whom they speak is in general neither identical to, nor a direct conceptualization of, the actual audience, the potential audience, or even the intended audience. For example, recall the story of Stokely Carmichael recounted in the previous section. Carmichael navigated the collapsed contexts and invisible audience by envisioning his audience as his Black activist supporters, though he knew that his radio and television messages would be heard by many others, including the white political leaders.

Until the advent of social media, the average person was not in the position of someone on broadcast media. But posting on social media puts even the participant in a casual conversation in the position of a public speaker, facing collapsed contexts and invisible audiences. There are empirical studies that support the view that social media users employ imagined audiences when posting (Marwick and boyd 2011, Litt and Hargittai 2016), and furthermore that they imagine different audiences for different posts, even when the potential audience remains the same (Litt and Hargittai 2016). To say that a speaker on social media imagines their audience is not to claim that speakers have *no idea* who will read a specific post (nor is this the case for broadcast media). Many speakers may have a very good sense of who will engage with a specific post given who has engaged with similar posts in the past or based on offline interactions; nevertheless, they are not in conversation with these people when posting, nor can they know that no one else will read or otherwise engage with the post. Their expectations of who will likely read or engage inform their imagined audience (Litt 2012). Litt further argues that factors like the norms and affordances of certain platforms, as well as features of the speakers themselves, such as what their motivations are in posting, all affect the imagined audience. Speakers sometimes report their imagined audience in general, abstract terms (e.g. "everyone", "the general public") and sometimes in specific terms (e.g. "co-workers", "people living in Baltimore", "friends who like reality TV"), with the same speaker varying both the specific conceptions and between general and specific for different posts (Marwick and boyd 2011, Litt and Hargittai 2016).[4] There is nothing in the concept of an imagined audience that prevents a speaker from imagining multiple audiences for a single post, though Litt and Hargittai noted that the participants in their study rarely had heterogenous specific imagined audiences in mind (such as close friends and co-workers). But this was only one study, and perhaps there are many speakers on social media who are carefully crafting message for multiple specific audiences at once, as politicians do with dogwhistles. In what follows I will write, for simplicity, about the speaker's imagined audience, but nothing I say is meant to exclude the possibility that the speaker is imagining multiple audiences.

[4] Some more examples from Litt and Hargittai's study: Personal ties such as "family", "close friends"; communal ties based on hobbies, interests, location, or political and religious affiliations, such as "all who were interested in good cleaning tips", "the local art community", "everyone with kids or who works in a school", "Anyone based in Portland", "all the people I know who are pro-life", "Christians"; professional ties such as "my coworkers", "my classmates", "my radio show audience", "lab colleagues"; and what they call phantasmal ties (to people the speakers couldn't be talking to or didn't really have a relationship with) such as "my deceased father", "my dog", "Obama". (p.5-6)

I also want to distinguish, conceptually, the imagined audience from the intended audience. This is a distinction that, as far as I know, is not explicitly developed in existing discussions of the imagined audience. I am thinking of the intended audience, roughly, as who the speaker intends to be the hearers of their speech act, who it is directed to, whose mental states they aim to act on, in the Gricean sense, while the imagined audience is the speaker's *conceptualization* of who the speaker is talking to, the way in which they are *envisioning* who they are talking to. In many cases, the imagined audience just is a conceptualization of the intended audience; they line up exactly. But they are distinct notions, and they do not always line up. In some cases, the imagined audience simply cannot be the intended audience because the speakers cannot have the intention that members of this audience hear the speech act or that their mental states be influenced, in other cases, the intended audience goes beyond the imagined audience. An example of the former are the people in Litt and Hargittai's study who said their dog or their deceased relative was their imagined audience.[5] The Carmichael case is a case of the latter: it is not that Carmichael does not *intend* for the white political leaders to hear or even to accept his message. They were among the people he had been talking to about civil rights. But in delivering his messages in a rolling pastoral tone, they are not his imagined audience.

As mentioned in the previous section, media scholars, including those writing on the topic of social media, are interested in how users navigate social norms and the presentation of self. When presentation of self is the topic at hand, the imagined audience is helpful in explaining what a speaker decides to say and how they say it. After all, one is forced to choose to present oneself in *some* manner or another and talk about *some* topic or another when faced with collapsed contexts and invisible audiences, whether that is in a way that is socially appropriate for one's childhood friends or one's work colleagues or some hybrid style. But one might object that there is no similar motivation for positing an imagined audience when it comes to the common ground. While it is true that a conversation with childhood friends would begin with a different common ground from one with work colleagues, there is no pressure to choose between these two (potential) common grounds when in a conversation with *both* childhood friends and colleagues. In fact, in a face-to-face conversation with childhood friends and colleagues, it would be decidedly odd (and rude) to presuppose information only shared with one's friends (or only with one's colleagues). Furthermore, there may not be social norms available to follow that are appropriate for all the participants in a conversation that involves a collapsed context (one might act too formally for one's childhood friends, or too informally for one's colleagues, or attempt some middle ground that doesn't feel quite right for anyone). However, there is *always* an appropriate common ground – the one that contains the actual presuppositions of the conversational participants, not including any information that is only available to friends or only to colleagues.

[5] In some of these cases, we might think that the imagined audience and intended audience very explicitly come apart in the mind of the speaker. For example, someone might write a post addressed to their baby (imagined audience) with the intention that their friends or family read their loving tribute to their baby on her first birthday (intended audience).

Transferring the lessons of this comparison to social media, one might argue that while we need the imagined audience to explain how speakers choose to present themselves, the common ground should simply be defined as the presuppositions held between the speaker and their entire potential audience, which may be the empty common ground (or close to it). This presents no real challenge from a practical or theoretical point of view. In practice, we have conversations with strangers with whom we share few presuppositions. And in theorizing about the dynamics of common ground, it is often simpler to assume we start with an empty common ground.[6]

There are both empirical and theoretical reasons for positing a common ground between the social media user and an imagined audience. The empirical evidence about imagined audiences cited seems to affect not just what social norms the users follow, but what information they presuppose, vocabulary they choose, and how they think their audience will interpret them (both in terms of illocutionary force and perlocutionary effect). As we will see in the next section, there is furthermore evidence from looking at social media posts themselves that suggests many of the authors presuppose a rich common ground. This is unsurprising, as it is often difficult and undesirable to communicate something that presupposes no common ground.[7] Social media users often communicate rich contents in short posts. Successfully communicating the content[8], force, and tone of a speech act, as well as the intended perlocutionary effects, all depend at least in part on a common ground. Posts on social media are also not all isolated from an explicitly linguistic context. They occur in threads and as replies that explicitly presuppose earlier posts in the same thread or posts to which they are replying.

3.2 Imagined Common Ground

In the absence of an actual audience, I propose that the speaker who posts on social media does so against a common ground with an imagined audience, whether that imagined audience is tacit or explicit. But what does it mean to have a common ground with an imagined audience (for short, *imagined common ground*)? Intuitively, the idea is that the imagined common ground is whatever propositions the speaker can take for granted when talking to their imagined audience, i.e. the things that would be mutually assumed were the imagined audience actual. For example, if someone is posting with their high school friends as the imagined audience, they can take for granted all sorts of facts about their high school, their home town, memorable high school experiences, etc. This suggests that the basic notion of imagined common ground is defined counterfactually: ϕ is common ground between a speaker and an

[6] It might be a bigger theoretical problem, however, that it would remain empty since *every* post in a conversation, not just the first one, is presented to an invisible audience with collapsed contexts.
[7] Goldberg (2021) makes related points in arguing against what he calls the maximalist position, which is taking the entire potential audience to be part of the exchange (and thus the basis for the common ground).
[8] Successfully communicating the referent of proper nouns, ambiguous terms, the content of context-sensitive expressions, as well as successfully communicating by employing certain expressions (such as specialized vocabulary) is all aided by the common ground, regardless of one's view on whether the common ground contributes to the semantic or pragmatic content.

imagined audience if ϕ would be part of the common ground were the members of the imagined audience actual conversational participants.

The question is how this could play the role it is supposed to play in message design (for the speaker) and in interpretation (for the hearer). As a starting point, we can think of it in terms of a Stalnakerian framework:

> It is imagined common ground that ϕ if the speaker *accepts* (for the purpose of the conversation) that ϕ, and acts as though their imagined audience accepts that ϕ, and acts as though all believe that all accept that ϕ, and acts as though all believe that all believe that all accept that ϕ, etc.

This captures the idea that the speaker acts as though certain propositions are taken for granted in writing their post. But since writing on social media is not in general an act of mere pretense, and audience members have to be in on the common ground to use it in interpretation, this doesn't fully explain its role in either message design or interpretation. The common ground has to be in some sense publicly available. An audience member must somehow (tacitly or explicitly) be able to recognize what is being taken as common ground in order for it to be useful in interpretation, and a speaker has to expect this of the audience in order to reasonably use it designing their message. We could add a clause to the Stalnaker-style definition that stipulates that ϕ has to be (in some sense to be spelled out) *recognizable* as common ground to a member of the imagined audience. Or, we might take this to be a good reason to explore whether a different framework for common ground is more illuminating.[9] Given space limitations, I leave this theorizing to future work, and aim only to show that the concept of an imagined common ground is coherent and can play the background role.

An example from danah boyd (2014)'s research on teenage use of social media may be helpful to illustrate what I have in mind. The following is from an interview with Hunter, a self-proclaimed "geeky, black fourteen-year-old" who "lives in two discrete worlds" (p.34) consisting of his cousins and sister on the one hand and the academically-minded "geeks" from his magnet school on the other:

> When I'm talking to my friends on Facebook or I put up a status, something I hate is when people who I'm not addressing in my statuses comment on my statuses. In [my old school], people always used to call me nerdy and that I was the least black black person that they've ever met, some people say that and I said on Facebook, 'Should I take offense to the fact that somebody put the ringtone, 'White and Nerdy' for me?' And it was a joke. I guess we were talking about it in school, and [my sister] comes

[9] For example, it may be promising to adapt Herb Clark's notion of a *shared basis* as the foundation of a common ground (Clark and Marshall 1981, Clark 1996).

> out of nowhere, 'Aw, baby bro,' and I'm like, 'No, don't say that, I wasn't talking to you.' (boyd 2014, p.34)

Hunter has an explicit imagined audience – his friends from school. According to the view I have proposed, the imagined common ground contains the information that he and his school friends share, and Hunter writes the post as though this background information can be taken for granted. This is not only an example of how speakers on social media do in fact assume specific common grounds (rather than what is in common with their whole potential audience), but also demonstrates how that common ground aids in communicative success – here it acts as necessary background for understanding that the post is a joke. How does the imagined common ground play its role in interpretation? In this case, there are clues in the content and the things referenced in the post that any of his friends reading it will recognize because, as he states, they were talking about it at school. He is presupposing *inside information*[10] with his school friends, and an insider would recognize it as such. What happens when someone outside the imagined audience reads Hunter's post, like his sister? In this case, she does not realize she is not part of the imagined audience, and since she does not actually have the shared background, she misses the point. She interprets Hunter as being sincere.

Conversations online reverse the usual order of things. The usual order of explanation is that *first* one knows who the conversational participants are, and consequently one knows the common ground. For example, it is first established that I am talking to a colleague, or fellow Montrealers, or friend from college, and as a result we can all presuppose (inside) information about work, or Montreal, or our college. In conversations online, a speaker *acts as though they are already* in a conversation with colleagues, or Montrealers, or friends from college by presupposing inside information, and when all goes well, members of the imagined audience recognize themselves as the imagined conversational participants by recognizing that they have the relevant inside information. While (a non-defective) common ground offline (if Stalnaker is right) ensures that the presupposed information is *actually* available, on the proposed definition, common ground in conversations online merely approximates this. And this seems right, given the invisible, potentially vast audience at the time of posting. I suspect that one important part of the explanation of the difficulties with successful communication on social media is that the imagined common ground is not actually public information. The speaker is acting as though certain presupposed information is available because they believe that their imagined audience shares these presuppositions and can figure out that the speaker does as well. But at the end of the day, this is still an *acting as if*, and what is being presupposed may not be so clear or so available after all.

There are many ways in which speakers can and do make an imagined common ground publicly accessible. Chosen vocabulary is one indicator; for example, using technical terminology about chess can be indicative of an imagined audience of fellow chess players, or using vocabulary

---

[10] Clark (1996) defines *inside information* as "particular information that members of the community mutually assume is possessed by members of the community". (p.101)

specific to a regional dialect can be indicative of an imagined audience of people from that region. Hashtags are another way in which information about imagined common ground can be made public, since hashtags are often associated with a particular movement or background (Hanteer et al. 2018). On platforms in which users have consistent identities attached to a history of posting (like Facebook, Twitter, YouTube, TikTok, and Instagram), whether the person is well-known, who their usual following is, the sort of content they usually post, what they have written in their biography and similar information can be indicative of the common ground.

In the following Instagram post by @KtownforAll, "an all-volunteer homeless outreach and advocacy organization in Koreatown, Los Angeles"[11], they clearly imagine their audience as other people located in Koreatown, Los Angeles who care about helping the unhoused.

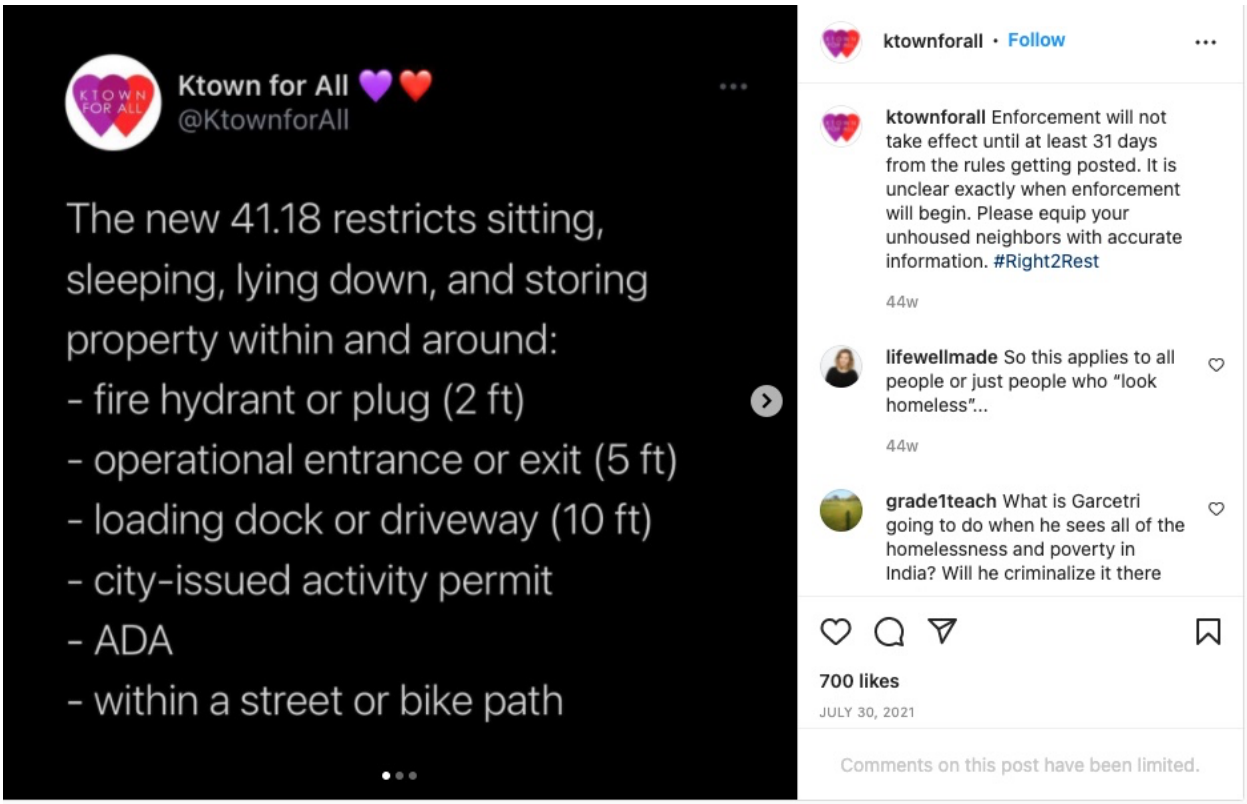

Example 1 https://www.instagram.com/p/CR-P1zslvzg/[12]

For example, the post presupposes that the referent of "the new 41.18" (which refers to a new ordinance in the Los Angeles municipal code) is common ground. A cooperative speaker who envisions themselves talking to a wider audience would say something like "the new ordinance of Los Angeles municipal code 41.18". In their comment alongside the post, @KtownforAll writes "Please equip your unhoused neighbors with accurate information". Here they explicitly address the audience, and they do so without qualification, even though they are talking about

[11] ktownforall bio, https://www.instagram.com/ktownforall/, last accessed Oct. 26, 2022
[12] Last accessed July 28, 2022

unhoused neighbors *of people in Los Angeles*. But this can be left unsaid, of course, given the right common ground. Finally, they use the hashtag "#Right2Rest" which seems to be a hashtag used to indicate content related to advocacy for the unhoused in Los Angeles, indicating clearly who they take their imagined audience to be, and a clue that would be recognized by such advocates.

What happens when someone outside the imagined audience reads the post? It is not uninterpretable – one can use clues from the wider context (e.g. the time and place it was posted, available facts about @KtownforAll, etc.) to figure out what they are referring to. In general, one of two things can happen when someone who is not part of the imagined audience reads a post. A reader might recognize they are (possibly) missing information for correct interpretation. They could react by not caring, and interpreting it as they see fit, or asking for clarification, or trying to fill in the missing information using other contextual clues. This kind of interpretive work can be more or less successful depending on the obscurity of the post, the extent of the context-sensitivity of the post, the availability of contextual clues, and the willingness and ability of the reader to do the work. Or, a reader might not recognize that they are not part of the imagined audience. In this case there will be a defective context, that is, what the conversational participants take to be the common ground doesn't line up.

Some posts presuppose so much common ground, they are simply uninterpretable without either already being in on it or coming to learn the relevant background information, for example, consider this joke:

Example 2 https://twitter.com/caseyjohnston/status/1166178548111085568?s=20[13]

This post presupposes as common ground specific information about Bret Stephens and going to his house to call him a bedbug that is provided in the two posts to which it is replying. It further presupposes knowledge of a joke proposal on Facebook in July 2019, to which over two million people replied, to storm Area 51 by having thousands or even millions of people Naruto run into Area 51 to see what the government is hiding there. Given the timing of their post (in August 2019), @caseyjohnston is reasonable in pre-supposing this information, at least among a certain group of social media users.

[13] Last accessed October 26, 2022. This tweet was the subject of the segment "Yes, yes, no" on the podcast Reply All #148 "Bedbugs and Aliens".

The fact that certain information is being taken for *granted* (even though it is at least sometimes in principle retrievable) becomes even more important when we look at different kinds of posts on social media. For example, mimetic posts that are meant to be satirical are only successful *qua* satire if the audience shares the relevant background information with the speaker. For example, when Richard Dawkins tweeted the following, other twitter users started posting their own tweets that began with "Bin Laden has won, in airports of the world every day" and continued with a non-consequential TSA inconvenience, without actually referencing or linking to the Dawkins tweet.[14]

Example 3 https://twitter.com/RichardDawkins/status/396956105869250561[15]

Someone who is not in the imagined audience might not even realize here that they are missing something, and may understand the speaker as sincere. The stakes are even higher when the same sort of satire is performed on a tweet that is racist, sexist, classist, etc. where the speaker relies on the common ground to be interpreted as making fun of the tweet rather than being racist etc. themselves.

The reverse holds true as well. Here is an example of something that is sincere (or at least, all evidence points to being sincere), but can easily be misinterpreted as satire if one does not presuppose the relevant common ground (perhaps something like *that the United States is a meritocracy*.)[16]

[14] Thank you to Nikki Ernst for pointing me to this example and for suggesting using the example of mimetic posts.
[15] Last accessed March 15, 2023
[16] In fact, that it appears to be satire was explicitly expressed in at least one reply:

Example 4 https://twitter.com/SeanODowd15/status/1403727466003369993[17]

Since audiences online are imagined and many people outside the imagined audience can and do read posts, misinterpretation is likely to occur more often than offline. How distant the misinterpretation is from the intended interpretation will depend on various factors, such as how far away from the imagined audience the reader is, how much time and effort they spend on understanding, to what degree the post itself is context-sensitive, and how many clues about the context are in the post itself. Furthermore, misinterpretations, and along with them, defective contexts, are more likely to persist online than offline.[18] Offline conversations commonly present opportunity to ask for clarification. A hearer can interrupt the speaker or immediately ask for clarification when the speaker finishes their sentence and a speaker can clarify what they meant when a hearer’s misinterpretation becomes obvious in the conversation. Furthermore, requesting and giving clarification is all perfectly common offline conversational behavior. Things look different online. A member of the audience cannot interrupt a speaker to request clarification, and because the conversational participants are not participating in a synchronous conversation, it takes more time for misinterpretation errors to be caught and addressed. Furthermore, there is no guarantee that the audience members who misinterpreted the speaker’s post will see the speaker’s clarification, nor is there a guarantee that the speaker’s clarification will reach the same audience as the reply that misinterprets it. In extreme cases, misinterpretations go viral but the clarifications that aim to address them reach far fewer people.

[17] Last accessed June 8, 2022

[18] Though see Peet (2023) for reasons why defective contexts might persist offline more often than Stalnaker thinks.

4 The dynamics of common ground

I have argued that for any given post, there is an imagined common ground that guides the speaker in deciding what to say and how to say it and the audience in interpreting it. If everything I have said so far is correct, then it establishes a role for common ground as the background for interpretation, that is, it establishes that there is at least a role for the common ground as a (series of distinct) static context(s). But what about the dynamic role of common ground, as the thing that is updated by speech acts and evolves over the course of the conversation? I want to put aside two features of the Stalnakerian picture that are not essential to a dynamic picture of context that includes a common ground. I will neither argue that the essential aim of a speech act is to change the common ground, nor that the central kind of conversation online is one that aims at inquiry. There are good reasons for rejecting these theses in general[19], but they seem particularly ill-suited for conversations online, where the central aim of many speech acts may be expressive (Lynch 2019, Goldberg 2021) or to gain 'points' (e.g. likes, followers, retweets, etc.) (Nguyen 2021). What I will argue is that the common ground tracks information added over the course of an exchange, and that the updated common ground plays a role in downstream interpretation and in constraining possible conversational moves. In general, we can make sense of updating the common ground in response to speech acts without accepting that speech acts are primarily characterized in terms of their effects on the common ground, nor accepting that the central goal of a conversation is inquiry. For example, suppose Alice tells Carol that her dog is sick because she wants sympathy and, furthermore, Carol doesn't care at all about finding out about Alice's pets, so neither party is interested in inquiry. Still, so long as Carol takes Alice to be sincere, or is at least willing to go along with her claim for the purposes of the conversation, it becomes shared background that Alice has a sick dog. This still constrains interpretation and appropriate moves. It would be inappropriate for Carol to say something that presupposes that she has no idea whether Alice has a dog.

One the most interesting differences between traditional public speaking and broadcast media on the one hand and social media on the other is that when it comes to conversations online, audience members can spontaneously interact with a speaker and with each other. In more traditional public speaking, there are limited opportunities for audience members to respond to a speaker. In many cases, applying the term 'conversation' to a politician delivering a televised speech is a bit of a strain on the term. However, social media is part public speaking and part private conversation: speakers put things out there for large, varied, and invisible audiences (like the public speaker) but at the same time engage in a back and forth with a limited number of people (like in ordinary conversation). In this respect conversations online resemble private conversation that take place in public settings.[20]

[19] See, for example, Dan Harris (2020) for arguments against characterizing assertion as aiming to update the common ground.

[20] But they are by no means exactly analogous, because while there are strong social norms against eavesdroppers joining a face-to-face conversation, the same norms do not apply online. People can read a conversation and join

I will argue that when a post is engaged with by another person, three things can happen (I am not claiming they *inevitably* happen). First, the initial imagined common ground (the input context) becomes actual. To become actual means that the common ground is between the speaker and an actual audience, rather than a merely imagined audience. This is not to say that the conversational participants necessarily envision themselves or their audience in the same way the speaker did, but merely to say that they hold the relevant presuppositions. Second, the content of the post has its usual effect on the actual common ground (for example, the proposition asserted is added). Third, a new imagined common ground comes into existence.

As noted in section 2.2, the main challenges for a theory of the dynamics of common ground is that online speech is not a manifest event and that mutual acceptance or belief is often difficult to achieve. Again, in the Stalnakerian tradition, it follows from the fact that a speech act is a manifest event that it automatically becomes common ground that it was made (the commonplace effect), and that if the semantics of the language is common ground and no one objects the content also becomes common ground (the essential effect). Similarly, presuppositions that go unchallenged are accommodated if they are not already part of the common ground. It is beyond the scope of this paper to examine whether this is the wrong picture in general (though I suspect it is), but this clearly won't work for conversations in which speech acts are not manifest events, including conversations online. In online conversations, assertions and presuppositions may go unchallenged because no one has seen them, they have been ignored, or they have been completely misunderstood.

Cases in which speech acts are not manifest events, such as conversations online, provide a novel reason for why we need *positive evidence* for something to become part of the common ground, not simply negative evidence as Stalnaker requires.[21] Herb Clark and others (Clark and Brennan 1991, Clark 1996) argue that, in general, we need positive evidence that our actions have been accomplished (such as the light on the elevator button indicating that we have been successful in calling the elevator). This is also true for our communicative actions, which Clark argues require positive evidence that they have been noticed and understood. He calls such evidence *grounding signals*: "To ground a thing… is to establish it as part of common ground well enough for current purposes." (Clark 1996, p.221). In writing about face-to-face conversation, Clark identifies four (non-mutually exclusive) main categories of grounding signals:[22]

1. Assertions of understanding: e.g. "uh-huh", "I see", "mmmm", nodding, smiling
2. Presuppositions of understanding: e.g. Taking a relevant next turn

---

at a later time. This is not to say that there are *no* norms about who can join a conversation on social media, just that the norms are different.

[21] An offline example of a conversation in which speech acts are not manifest events is when two people are talking to each other from different rooms in the house. If the hearer does not reply to the speaker, it is actually odd for the speaker to assume that the hearer accepted what they said rather than didn't hear what they said.

[22] Clark & Brennan 1991 note that different media will have different grounding signals.

3. Displays of understanding: e.g. An answer displays (in part) how someone construed a question
4. Exemplifications of understanding: e.g. Offering a paraphrase or verbatim repetition, grimacing, looking disappointed, etc.

When somebody makes an appropriate reply to a post – that is, one that demonstrates understanding of the post – this provides positive evidence that the post was seen and understood (and in some cases, agreed with). In Clark's terms, appropriate replies to posts can be displays of understanding, since they can display how the person who replied construed the original post. They can also be presuppositions of understanding, since an appropriate reply is likely to be a conversational move that best makes sense given the assumption that the original post has been understood. *Reactions* provided by certain platforms (such as liking, loving, upvoting) are also (weaker) grounding signals, but I will here focus on replies because reactions don't prompt a subsequent conversational move.[23] Positive evidence can help overcome the lack of manifestness. When someone makes an appropriate reply, there is evidence that there are now two actual conversational participants who presuppose the relevant propositions in the initial, input imagined common ground. Furthermore, if the reply indicates acceptance of the content, then this is reason to take that content also as actually shared between the conversational participants.

Consider the following tweet and two distinct replies to it in examples (5) and (6). The first exemplifies someone who accepts the input imagined common ground, but rejects the assertion, while the second accepts both the input common ground and the assertion.

[23] Clark notes that some grounding signals provide stronger evidence than others; e.g. for his categories, 1 and 4 are weaker than 2 and 3.


Tweet
Tim Pool
@Timcast
Johnny Depp's win is a huge blow to "believe all women" and corporate press credibility
6:08 PM · Jun 1, 2022 from Maryland, USA · Twitter for Android
1,901 Retweets
94 Quote Tweets
22.1K Likes
The narcissism of small difference @Non_Ideologue · 16h
Replying to @Timcast
Nah, they'll just double down.
7

Example 5

Example 6

In (5), @Non_Ideologue explicitly rejects @Timcast's assertion, but accepts the presuppositions of the input imagined common ground (that Johnny Depp won the defamation trial v. Amber Heard, the presuppositions associated with calling the "believe women" movement "believe all women"[24], etc.). In fact, it is necessary to take those presuppositions into account to correctly interpret the reply – that 'they' refers to those who support the so-called "believe all women" movement as well as the corporate press. In example (6), @Mannyquintero's reply is evidence of accepting the relevant presuppositions in the imagined common ground as well as the proposition asserted in @Timcast's post. In terms of grounding signals, his post includes explicit agreement – "yeah, the tide is shifting" as well as rewording of/expanding on the central point ("we shall see true equality now"). The common ground seems to play its usual role here both in downstream interpretation and constraining possible moves. When @Mannyquintero says "the tide is shifting" it is more or less clear what tides – the tides towards "true" equality of men and women and against the so-called "believe all women" movement. @Mannyquintero can trust that @Timcast (and potentially other conversational participants) will know what tides he is referring to. Furthermore, neither conversational participant can now go on to say something that contradicts the presupposition that Johnny Depp won the trial or its effect on the so-called "believe all women" movement without retracting what they said earlier.

[24] "Believe all women" is a strawman of a slogan, attributed by right-wing pundits and their followers to the left, to feminists, and to members of the #MeToo movement, who actually used the slogan "believe women". "Believe all women" implies that the people it is attributed to are arguing that every single woman should be believed all the time, an easy slogan to counterexample, but also one that nobody actually holds.

There are two important objections to address at this point. One is the question of whether positive evidence of understanding or acceptance is sufficient for achieving the iterated propositional attitudes of a typical Stalnakerian common ground. The second is whether a common ground really constrains conversational moves online in a way analogous to conversations offline. These are both important complications that require more detailed exploration in future work. Here I will only address them insofar as to show the theory I am proposing is viable.

It is right that there is no *guarantee* of iterated acceptance or belief – there is always the possibility that the original poster has not read the reply (yet), and moreover the replier knows this. But the positive evidence in a reply gets us close enough that it is fruitful to posit an actual common ground here. First, as I've already argued, it plays its traditional role in downstream interpretation and constraints on moves. For example, if @Timcast replies to @Mannyquintero he is constrained by the common ground as I have described it. Second, the common ground is often idealized in this way in accounting for other kinds of conversations that lack manifestness (and simultaneity) in some way, such as letter writing, emails, or text messages. When someone sends an email reply, until they receive a reply or other indication of receipt, there is always some doubt as to when or if the other person will read it.[25] While there might be some cases of conversations online in which the communicative channels are very noisy and we thus don't want to say common ground is achieved, just like there are cases of letter-writing and emailing that involve noisy communicative channels, it doesn't mean that we want to say there is no common ground in the good cases. Evidence has to end somewhere.

The second objection is that I have misdescribed things – conversations online work quite differently from offline conversations. People ignore conversational moves all the time online. There is no expectation that everyone replying to a post is engaging with each other or has read all the other replies or comments in a thread. So in what sense are those replying constrained by a common ground? It is true that there is not in general a common ground that constrains the moves of everyone who engages with a particular post. The evolution of the common ground more tightly constrains the interlocutors who have already engaged with each other than it does newcomers who engage with the original post or any replies to it. In many cases, it does not constrain the possible moves of newcomers at all, and it is best to understand what is happening as the beginning of a *new* conversation that happens to have an overlapping utterance (in some sense) with another conversation. For example, I think it's best to understand @Timcast and @Non_Ideologue in one conversation while @Timcast and @Mannyquintero are in a completely different one, since there is no interaction between @Non_Ideologue and @Mannyquintero. In other cases, the common ground does constrain the moves of newcomers, in particular, the cases in which the newcomers are directly engaging

[25] This is different from, but related to, what Harris (2020) calls a *publicity averse* speech act, which requires there to be a positive reason to think that common belief (or acceptance) cannot be achieved, such as communicating via messenger through an enemy-filled valley or emailing someone who often ignores their emails.

with all the interlocutors in an existing conversation.[26] That is to say, the individuation of conversations online looks different from offline conversations. While a single utterance offline is relatively rarely part of multiple conversations, a single post on social media is quite often part of multiple conversations. Given the norms of interaction online, it is perfectly appropriate to start a brand-new conversation by replying to something said in a different conversation. We can think of this as a sort of branching tree model of conversations, where each post is a node that can branch into possibly many conversations.

Finally, upon writing any post, whether it is an initial one or a reply or comment in a thread, a speaker is faced with the same invisible, diverse audience. Not only is a speaker faced with the *possibility* that potentially many people from diverse social groups will read their reply to a particular post, in many cases they *aim* to reach a new audience beyond the current conversational participants. This is unlike ordinary face-to-face conversations in more ways than one. As already established, a speaker in a face-to-face conversation is generally talking to the existing conversational participants and even if other people are in earshot, they do not count as the audience for the speech act. But moreover, it is rare for a speaker having a face-to-face conversation with others in earshot to intend for those outside the conversation to hear, be amused by, engage with, or otherwise react to their utterance.[27] However, online, this is the norm. Replies or comments directed at existing conversational participants are often equally directed at a new audience (as we saw very clearly, for example, in example (2), @caseyjohnston's post that made the joke about area 51 for going to Bret Stephen's house). This adds a serious complication to the dynamic picture of context: each move in a conversation doesn't simply add to an existing common ground, nor does it simply actualize and add to the input imagined common ground, but in addition adds an overlapping, separately evolving new imagined common ground. (Whether this results in an overlapping beginning to a new conversation or not is a distinct matter.)

5 Conclusion

I have argued that common ground plays an important role in a theory of context for conversations online. It helps explain how speakers communicate what they do, including rich contents conveyed by short posts, various illocutionary forces and intended perlocutionary effects (such as amusing, getting likes, enraging, etc.), and also manage to successfully tell jokes and perform acts of satire without giving explicit background. The picture of common ground is different than the traditional one given for offline conversations, and the differences also help explain the abundance of unsuccessful communication online. In section 3, I argued that the

[26] This brings up an interesting question of whether what is best understood as two separate conversations can at some point become one, for example, if someone came along and engaged with @Timcast's post in a way that replies to *both* @Non_Ideologue and @Mannyquintero. Whether this is viewed as bringing what was two conversations into one or starting a new conversation will likely depend on the details.

[27] It is not the norm but it is also of course not unheard of for someone to say something hoping someone within earshot will overhear.

imagined common ground is not public in the same sense as an actual common ground, rather it is an *acting as if* certain information is public, which can be more or less successful depending on the speaker and other circumstances. Furthermore, since posts are often read by people outside the imagined audience (who may or may not realize they are outside the imagined audience), there is further opportunity for misinterpretation even when the presuppositions in the common ground are obvious to those in the imagined audience. Finally, in section 4, I argued that the common ground does not evolve in a straightforward way because every conversation involves an actual common ground with overlapping imagined common grounds. Moreover, the boundaries of a conversation are unclear. Together, this means it is often unclear how exactly the common ground constrains subsequent conversational moves.

Works Cited

boyd, d. (2014). *It's complicated: The social lives of networked teens.* Yale University Press.

Clark, Herbert H. & Marshall, Catherine R. (1981). "Definite Reference and Mutual Knowledge". In Aravind K. Joshi, Bonnie L. Webber & Ivan A. Sag (eds.), *Elements of Discourse Understanding*. Cambridge, UK: Cambridge University Press. pp. 10–63.

Clark, Herbert H. & Brennan, Susan E. (1991). "Grounding in communication". In Lauren Resnick, Levine B., M. John, Stephanie Teasley & D. (eds.), *Perspectives on Socially Shared Cognition*. American Psychological Association. pp. 13--1991.

Clark, Herbert (1996). *Using Language*. Cambridge University Press.

Coddington, M, Lewis, A & Belair-Gagnon, V. (2021). "The Imagined Audience for News: Where Does a Journalist's Perception of the Audience Come From?" *Journalism Studies* 22(8), 1028-1046.

Goldberg, Sanford C. (2021). "The Promise and Pitfalls of Online 'Conversations'". *Royal Institute of Philosophy Supplement* 89:177-193.

Harris, Daniel W. (2020). "We talk to people, not contexts". *Philosophical Studies* 177 (9):2713-2733.

Lynch, Michael P. (2020). *Know-it-All Society: Truth and Arrogance in Political Culture.* New York, NY, USA: WW Norton.

Marwick, A. E., & boyd, d. (2011). "I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience". *New Media & Society, 13*(1), 114–133

Meyrowitz, J. (1985). *No sense of place: The impact of electronic media on social behavior.* Oxford University Press.

Nelson, J.L. (2021). *Imagined audiences: How journalists perceive and pursue the public.* Oxford University Press.

Nguyen, C. Thi (2021). "How Twitter gamifies communication". In Jennifer Lackey (ed.), *Applied Epistemology*. Oxford University Press. pp. 410-436.

Ong, W.J. (1975). "The writer's audience is always a fiction." *PMLA/Publications of the Modern Language Association of America*, 90(1), 9-21.

Peet, Andrew (2023). "Defective Contexts". In Rachel Katharine Sterken & Justin Khoo (eds.), *Routledge Handbook of Social and Political Philosophy of Language*. Routledge.

Stalnaker, Robert (1978). "Assertion". *Syntax and Semantics* (New York Academic Press) 9:315-332.

Stalnaker, Robert C. (1999). *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford, GB: Oxford University Press UK.

Stalnaker, Robert (2002). "Common ground". *Linguistics and Philosophy* 25 (5-6):701-721.

Stalnaker, Robert (2014). *Context*. Oxford: Oxford University Press.