

# Common Ground: From Formal Pragmatics to Psycholinguistics

Paula Rubio-Fernandez and Daniel W. Harris<sup>\*</sup>

April 29, 2025

## Abstract

Common ground is the information that the participants in a conversation treat as background information for the purposes of their interaction. We review two traditions of research on common ground: The formal tradition, consisting mainly of theoretical linguists and philosophers of language, has developed increasingly sophisticated formal models of common ground in order to generate predictions about an expanding range of empirical phenomena. Meanwhile, the psycholinguistic tradition has focused on a narrower range of phenomena while developing more realistic theories of the psychological mechanisms that allow us to select and represent common ground. After summarizing these two traditions, we consider several reasons why they should be re-integrated, and argue that the best way to bring them back together would be to adopt a *cognitive-pluralist approach*, whereby language users have access to a variety of mechanisms for managing background information, which are more or less available and efficient depending on the communicative situation and the kind of information mentally represented, as well as the cognitive demands of each mechanism.

## 1 INTRODUCING COMMON GROUND

Common ground is the information that two or more interlocutors treat as shared background information for the purposes of an interaction. We can illustrate this idea with a minimal pair of examples (cf. Heal 1978; Lederman 2017). Suppose

---

<sup>\*</sup>(acknowledgments)

that two friends, Sam and Ethan, have spent election night together, avoiding news coverage. Then, one of two things happens:

**TV** While Sam and Ethan are sitting on the sofa together, Ethan turns on the TV, whereupon he and Sam together see an image of the incumbent celebrating victory.

**PHONES** Sam and Ethan separately, and coincidentally receive text messages from friends that consist in the same image of the incumbent celebrating victory.

In both scenarios, Sam and Ethan each come to believe that the incumbent has won. However, in **TV**, this information also takes on an additional, *public* status between them that it lacks in **PHONES**: it becomes common ground.

This status has been argued to play a wide range of roles in conversation, shaping both the design and interpretation of linguistic utterances. For example, common-ground theorists have analyzed assertions as proposals to update the common ground, and argued that it is therefore often infelicitous to assert what is already common ground (Stalnaker, 1978, 325). This explains why it is strange for Sam to utter (1) in **TV** but not in **PHONES**:

- (1) I have just seen the incumbent, victorious. #TV/✓Phones

Another influential idea is that, when an interlocutor makes an utterance with a presupposition, this presupposition must be common ground. Or, at least, interlocutors must be willing and able to accommodate it—i.e., to immediately identify what was presupposed and proceed as though it was already common ground (Lewis, 1979). Our minimal pair bears this out: When used as an attitude verb, ‘surprise’ triggers a presupposition that its compliment is true, which is why (2) is strange in **PHONES** (where it is not common ground that both Sam and Ethan have seen the incumbent celebrating) but not in **TV** (where the information is public):

- (2) Are you surprised to see that the incumbent has won? ✓TV/#Phones

The roles of common ground in assertion and presupposition combine to explain the following minimal pair (adapted from Lewis 1979):

- (3) a. The incumbent has children, and all the incumbent’s children are excited.  
b. # All the incumbent’s children are excited, and the incumbent has children.

By uttering the first clause of (3a), a speaker proposes to add the proposition that the incumbent has children to the common ground, and this, in turn, satisfies the presupposition that is triggered by ‘All the incumbent’s children’ in the second clause. By contrast, with the clauses’ order reversed in (3b), the presupposition that the incumbent has children needs to be accommodated in order to interpret the first clause, making the second clause redundant, and so, infelicitous.

Common ground also plays a special role in the use and interpretation of context-sensitive expressions. For example, quantifier phrases, such as ‘everyone,’ are normally contextually restricted in a way that seems to be guided by common ground. We can see this from the fact that it is much easier to imagine a felicitous utterance of (4) in TV than in PHONES, presumably because the extra information in common ground would make it easier for Sam and Ethan to coordinate on an appropriate restriction for ‘everyone’ (e.g., supporters of the incumbent).

(4) Everyone is going to be excited. ✓TV/#Phones

Similarly, the acceptability of (5) in TV but not in PHONES seems to arise from the fact that the information in the common ground plays a special role in the design and interpretation of pronouns and other definite noun phrases: Roughly, one should utter a definite noun phrase only if the common ground will make it obvious who the referent is (Clark and Marshall, 1981; Stalnaker, 1978):

(5) He looks excited. ✓TV/#Phones

The dual roles of common ground in recording the effects of upstream speech acts and in determining the felicity and interpretation of downstream speech acts makes it useful for explaining anaphoric phenomena, in which the meanings of later utterances depend on those of earlier utterances. For example, in (6), the pronoun, ‘he,’ is most naturally understood as anaphoric on “the incumbent:”

(6) I am watching the incumbent celebrate. He looks excited!

One promise of common-ground theories is to offer a unified explanation of anaphoric uses of pronouns (as in (6)) and deictic uses (as in (5)). In both scenarios, the sentence, “He looks excited,” updates the common ground with information about whichever male it already represented as most prominent. But whereas this prominence has extralinguistic causes in (5), it is owed to the pronoun’s antecedent in (6).

Following this informal characterization of common ground, we will loosely organize our review around a historical narrative: In the 1970s, a number of philosophers and linguists developed an influential formal-pragmatic model, with a representation of common ground at its center. This model became influential, in

part because it offered a unified account of a wide variety of pragmatic phenomena. However, the model was also idealized in a number of ways, abstracting away from both a variety of other pragmatic phenomena and the details of the underlying psychological mechanisms. Much subsequent work on common ground can be understood as a series of attempts to lift these idealizations, but this work largely split into what we will call the *formal* and *psycholinguistic* traditions. Whereas the formalists have complicated the model to cover a wider range of conversational phenomena, the psycholinguists have mostly focused on one phenomenon (definite reference) while seeking more realistic accounts of underlying cognitive mechanisms.

Each of these traditions has made significant progress toward their respective aims, but have become less integrated over time. In §2, we will summarize key developments from the two traditions. In §3, we will argue that the two traditions' insights are more compatible and mutually illuminating than has been supposed. Finally, in §4, we will argue that the best way to combine the two traditions is by means of a pluralist approach to the cognitive mechanisms underlying common ground.

## 2 HISTORICAL OVERVIEW

A theory of common ground first took shape as part of a formal-pragmatic model developed by Stalnaker (1970; 1973; 1974; 1978) and Karttunen (1974). This model is built using the tools of possible-worlds semantics, in which a proposition is modeled as the set of possibilities (or *possible worlds*) in which it is true (see Figure 1). Common ground is then modeled as a set of propositions. Derivatively, the *context set* is the set of possible worlds that are compatible with all of the propositions in the common ground—i.e., the intersection of all of the propositions in common ground. In effect, the context set is the set of possibilities that the interlocutors are treating as live candidates for actuality, and which they are seeking to decide between in conversation. Conversations are modeled as cooperative exchanges in which participants take turns adding propositions to the common ground by asserting them. By adding new information to the common ground, assertions eliminate possibilities from the context set, thereby whittling down the candidates for actuality. By updating the common ground in this way, interlocutors also shape the future of the conversation by constraining which assertions and presuppositions are felicitous, changing which uses of anaphoric and context-sensitive expressions are available, and so on. Stalnaker's model is therefore a *dynamic* model,




Figure 1: The common-ground-based model of conversation developed by Stalnaker, Karttunen, and others.

in the sense that it treats common ground as an evolving body of information that is altered by speech acts in ways that change which speech acts are subsequently performable.

What is this model a model of? The original answer, and the one that remains most influential in the formal tradition, is that common ground is the body of propositions to which the interlocutors bear the right kind of *iterated propositional attitudes*.<sup>1</sup> This line of thought originally grew out of earlier work on *common knowledge* (a.k.a. *mutual knowledge*; we use these terms interchangeably) (Lewis, 1969; Schelling, 1960; Schiffer, 1972). The most influential definition of common knowledge invokes intersubjectively iterated propositional attitudes, as follows:

- (7) **Common Knowledge (iterated)**  
 A and B commonly know that  $p$ , if and only if (ia) A knows that  $p$ ; (ib) B knows that  $p$ ; (iia) A knows that B knows that  $p$ ; (iib) B knows that A knows that  $p$ ; ...

One possible treatment of common ground is to identify it with the set of propositions that the interlocutors commonly know, in the sense of (7). This would explain the difference between our two scenarios from §1: In (PHONES), Sam and

<sup>1</sup>A propositional attitude is any mental representation that involves an agent taking an attitude (e.g., belief, knowledge, desire, hope, intention) to a proposition (e.g., that *the incumbent has won the election*).

Ethan each know who won the election (satisfying clauses (ia–b) of (7)), but neither knows that the other knows, or knows that the other knows that they know, etc., and it is their possession of this further knowledge that constitutes the information’s publicity in TV. A number of scholars have argued that satisfying (iia–b) is also insufficient, however: If Sam knows that Ethan saw the election result but thinks that he thinks that she didn’t see it, she might not treat this information as common ground (cf. Schiffer 1972; Clark & Marshall 1981). Increasingly elaborate thought experiments of this kind led to infinitely iterated definition, like (7).

However, although some theorists who define common ground in terms of an iterated attitude would say that the attitude in question is *by default* knowledge (Yalcin, 2024), nearly all would deny that it is *always* knowledge. People sometimes unwittingly treat falsehoods as common ground, for example, which suggests that a more general definition of common ground would identify it with *common belief*, which can be defined by replacing each instance of ‘knows’ (7) with ‘believes.’ Even more influentially, Stalnaker (2014) has argued that common ground is made up of the propositions that we commonly *accept for the purpose of the conversation*. “To accept a proposition,” Stalnaker (1984, 79) tells us, “is to treat it as a true proposition in one way or another—to ignore, for the moment at least, the possibility that it is false. One may do this for different reasons, more or less tentatively, more or less self-consciously, with more or less justification, and with more or less feeling of commitment.” This view is motivated by cases in which what we believe comes apart from what we treat as common ground—e.g., when someone says something with which we disagree, but we let it slide in order to avoid conflict, when we engage in extended pretense or fictional discourse (Stokke, 2023), or when we communicate things via insinuation that we would not openly acknowledge (Camp, 2018).

We will use the term ‘iterated theories’ to refer to the broad family of views that identify common ground with the set of propositions that interlocutors take infinitely iterated attitudes toward. In §3.1, we will consider the objection that these theories are psychologically unrealistic. However, it is worth emphasizing that, while the formal-pragmatic model originally came packaged in terms of iterated propositional attitudes, the two ideas are separable: The formal model could be paired with some other theory of how agents represent and arrive at the shared information represented by the model.

As its creators acknowledge, this formal model of conversation “is highly idealized and oversimplified, best regarded as an artificial game designed to model some of the features of the structure of a discourse” (Stalnaker, 2014, 89). For example, although Stalnaker informally theorizes about *defective contexts* (i.e., situations where interlocutors treat conflicting information as common ground), his

formal model does not aim to make predictions about what happens in those situations (Stalnaker, 1978, 322). It also ignores speech acts that can't be thought of in terms of cooperative information sharing, such as questions or commands. And it tells us little about the cognitive mechanisms whereby we accumulate and maintain common ground. As we will see in §§2.1–2.2, much subsequent work on common ground has been concerned with finding ways to lift these idealizations, thereby turning the model into a more comprehensive and realistic representation of conversation.

Given how heavily idealized the common-ground model initially was, one might wonder why it became so influential. The answer is that it was used to understand and connect a range of pragmatic phenomena—including those that we discussed in §1—which had hitherto been approached as an array of unrelated puzzles. It is the model's unifying explanatory power that explains its initial and enduring influence.

## 2.1 THE FORMAL TRADITION

The formal tradition has been focused on explaining a wider variety of conversational phenomena, often by adding complications to the common-ground model. An early, influential example is David Lewis's (1979) model of *conversational score*, which, like the score of a baseball game, is a dynamic record of various kinds of information about what has happened so far in the "language game" that are relevant to determining what can happen next. Alongside a body of presupposed information (i.e., Stalnaker's common ground), Lewis posited six additional components of conversational score, each of which can be updated in characteristic ways by certain speech acts, and which change which speech acts are possible going forward.

One influential example is Lewis's idea of a *permissibility sphere*, which changes in response to authoritative speech acts of command and permission (the boss says "Stay late"), and which affects which subsequent statements about obligation and permission are true (the employee says "I have to stay late"). This was the first in a long tradition of proposals about how to model the updates and downstream consequences of the speech acts that we perform using imperative sentences (e.g., Portner 2007; Kaufmann 2012). A second example is Lewis's idea that the context includes shifting standards of precision for vague predicates (e.g., are we using "tall" to mean *tall for a basketball player* or *tall for a toddler*?). This inaugurated a tradition of work in which the common ground is used to model not just information about the extra-conversational world, but also information about the conversation itself, including participants' short-term agreements about how to

use context-sensitive expressions within it (e.g., Barker 2002; MacFarlane 2016). Lewis’s general strategy—positing complications to common ground in order to capture more phenomena—has been even more influential than his specific proposals about what to add.<sup>2</sup>

We now want to zoom in on two influential proposals about how to complicate common ground. Our first case study is the proposal that, in addition to a body of presupposed information, interlocutors must keep track of a shared representation of the question(s) that they are seeking to answer (Beaver et al., 2017; Carlson, 1982). For example, in Roberts’ (2012b) influential model, the Stalnakerian common ground is accompanied by the *question stack*, which is an ordered list of questions, the topmost of which is the *question under discussion* (or QUD). Just as asserting a proposition is proposing to add it to the common ground, to ask a question is to propose it as a new QUD. If Stalnaker’s context set represents the live possibilities that interlocutors are trying to choose between as the overarching aim of their “joint inquiry,” then the questions in Roberts’ stack embody “strategies of inquiry”—shared sub-goals to direct the inquiry toward more manageable and specific topics. She captures this idea by modeling questions as sets of mutually exclusive answers (following Hamblin 1973), and by thinking of the QUD as partitioning the possibilities in the context set into a collection of subsets, each of which represents a possible live answer to the QUD. The interlocutors’ immediate conversational goal is then to answer the QUD by making assertions that eliminate some of the cells of this partition.

Roberts uses this model to predict our judgments about which speech acts are relevant: An assertion is relevant only if it partially answers the QUD, which is to say that it rules out at least one live answer. A question is relevant if it poses a live subquestion of the QUD, which is a question not already answered by the common ground, each of whose complete answers would partially answer the QUD. Roberts uses this definition of relevance to refine Grice’s (1975) theory of relevance implicature. On Roberts’ view, the search for an indirect meaning is triggered when the speaker asserts something that, taken literally, is not relevant to the QUD (and

---

<sup>2</sup>Subsequent work has posited further complications to explain, e.g., non-assertoric speech acts, such as commands, questions, and exclamations (Farkas and Bruce, 2010a; Harris et al., 018b); conversational implicature, indirect speech acts, and insinuation (Asher and Lascarides, 2001; Camp, 2018); non-factual assertions and disagreements (Gibbard, 2003; Yalcin, 2012); lying and deception (Stokke, 2018); interlocutors’ strategies for resisting and negotiating speech acts (Bledin and Rawlins, 2018; Farkas and Bruce, 2010a); the difference between conversational contributions that can be directly responded to and those that are backgrounded (Simons et al., 2017); fictional discourse (Stokke, 2023), and hate speech (Langton, 2012; McGowan, 2019).



therefore not cooperative), and the search for indirect meanings is constrained by the assumption that the speaker must have implicated something that *is* relevant. Roberts also argues that the function of prosodic focus is closely related to QUD: Roughly speaking, when making an assertion, it is felicitous to focus an expression only if replacing that expression with an appropriate *wh*-phrase would yield a sentence whose content is the QUD (e.g.: “Who is here?” “JIM is here”/# “Tom is HERE”). On this view, the function of prosodic focus is to probe and reinforce the QUD.<sup>3</sup>

Although the question stack and QUD are additional elements of context, over and above the common ground of the original formal model, it should also be clear that the QUD model can’t be understood without common ground, and also gives the common ground important new jobs to do: The model generates predictions about relevance, implicatures, and prosodic focus as a function of *both* QUD *and* common ground. These kinds of interactions represent one kind of reason for a common-ground theorist to think about at least some common-ground-external elements of conversational score.

Our second case study is the idea that we should think of common ground as including object representations in addition to propositions. Lewis himself formalized this idea by including a salience-ranked list of entities among the elements in conversational score (Lewis 1979; cf. Stojnić 2021), but a much more influential strategy has been to rethink the representation of common ground itself, modeling it as (or as including) a collection of *discourse referents* (Karttunen, 1976) or *file cards* (Heim, 1982, 1983a), which are shared representations of the people and things under discussion. Another modeling strategy has been to modify Stalnaker’s context set, so that each element is not just a possible world, but a pair consisting of a possible world together with an assignment function, which is a possible mapping from noun phrases to their referents (Groenendijk and Stokhof, 1991; Mandelkern, 2024). Heim (1983b, 228–229) shows how to think of these sets of world–assignment pairs as abstract representations of sets of her file cards, much as Stalnaker’s context set is an abstraction from common ground. The intuitive idea behind these modeling approaches is that our assertions don’t just add information about the world to common ground, thereby eliminating possibilities from

---

<sup>3</sup>Others have used QUD models to predict data about, *inter alia*, epistemic modals (e.g., “might,” “probably”) (Dever and Schiller, 2021; Yalcin, 2007), propositional-attitude verbs (“knows,” “believes”) (Schaffer, 2007; Yalcin, 2018), contextual enrichment of utterances’ literal meanings (Schoubye and Stokke, 2016), loose speech (Hoek, 2018), disjunction (Simons, 2001), and presupposition projection (i.e., data about when sentences inherit the presuppositions of their parts) (Simons et al., 2017). This explanatory versatility has made QUD models highly influential.

the context set: Assertions also add information about the discourse itself—e.g., by introducing new people and things as topics for discussion, or making familiar things more salient. This is modeled as the introduction of new object representations to common ground, or as the manipulation of the assignment functions that are available in the context set. We will group these theories together as advocating for *objectual common ground* (i.e., common ground as including representations of objects), as opposed to Stalnaker’s propositional common ground (a set of propositions).

Objectual common ground is motivated as the central component of dynamic theories of anaphora, which interpret indefinite noun phrases (e.g., “a woman”) as devices for introducing new object representations to the common ground, and definite noun phrases (e.g., “she”, “the woman”) as devices for adding information to the object representations already available.<sup>4</sup> We won’t attempt a detailed summary of the motivations for these theories, many of which involve intricate patterns of interaction between grammatical phenomena and common ground. But here is a simple illustration: Unless it is common ground that a certain woman is particularly salient, it would normally be baffling to start a conversation by saying, “She is a banker.” We might explain this by saying that the common ground does not help us to identify a referent for “she.” However, one *can* begin a conversation by saying, “I know a woman. She is a banker.” Uttering the first sentence must therefore change the common ground in some way that licenses the second sentence. Propositional theories struggle to explain this, since it was already presumably common ground that the speaker knows a woman, and so their assertion couldn’t have added a new proposition to the common ground. By contrast, objectual theories can explain the contrast by saying that it is felicitous to utter a pronoun, such as “she,” only when there is a salient female-marked object representation in the common ground, and the function of uttering “I know a woman” is to create and make salient just such an object representation. It is crucial for this explanation that, as Heim (1982, 249) put it, “discourse referents are not referents”; rather, they are shared representations that connect anaphora to their antecedents, and whose relationship to actual entities in the world may be more or less underspecified.

We have tried to give a taste of the myriad ways in which formal models of common ground have been complicated, and why. However, these technical innovations and empirical advances raise questions about the underlying psychol-

---

<sup>4</sup>In dynamic *semantic* systems, these context-updating functions are built into sentence meanings, whereas dynamic *pragmatic* systems reduce context update to pragmatic reasoning (Lewis 2014; 2021).

ogy of conversational score—about what the formal models are modeling—which have mostly gone unaddressed. Lewis’s own remarks about how conversational score relates to interlocutors’ psychology are cryptic and noncommittal (Lewis, 1979, 346). Roberts sometimes says that the question stack represents interlocutors’ goals or plans (Roberts, 2012b, 3, 12, 26), but these remarks have not yet been developed into a detailed, empirically testable theory. And although Heim (1982, 288) presents her theory of common ground as an extension of Stalnaker’s, it seems that Stalnaker’s reduction of common ground to iterated propositional attitudes can’t be straightforwardly extended to objectual models, since object representations aren’t propositional. So, although the formal tradition has made progress on lifting some of the common-ground model’s idealizations—those dealing with its narrow empirical scope—it has simultaneously deepened the model’s psychological idealizations.

## 2.2 THE PSYCHOLINGUISTIC TRADITION

H. Clark and colleagues proposed the first psycholinguistic model of common ground (Clark, 1996; Clark and Carlson, 1981; Clark and Marshall, 1978) as a response to what they called the *processing paradox of mutual knowledge* (Clark and Marshall, 1981): While the formalists’ arguments for the need for mutual knowledge in human communication were impeccable (Lewis, 1969; Schiffer, 1972), mentally computing an infinite recursion of propositional attitudes would take an infinite amount of time—contrary to our everyday experience of conversation. To solve this processing paradox, Clark and Marshall (1978; 1981) argued that, instead of mentally representing common ground as iterated propositions, interlocutors must rely on *heuristics for assessing mutual knowledge*. In other words, interlocutors must find the right shared bases—or evidence, for assuming common ground.

According to Clark (1996), common ground is the foundation for all joint action, including the construction of speaker meaning and its understanding by the addressee. However, the proposed heuristics for assessing mutual knowledge have a much narrower scope, as they are formulated in relation to definite reference alone: When interlocutors make or interpret a definite reference, they look for evidence of so-called *triple co-presence*; namely, an event including the speaker, the addressee and the intended referent (Clark and Marshall, 1978, 1981). The strongest evidence of triple co-presence comes from physical co-presence (e.g., when two friends are jointly looking at a painting in a museum), followed by linguistic co-presence (e.g., when one of these two friends later refers to ‘the beautiful Picasso’). The third heuristic for assuming common ground is cultural co-

presence, which includes all the knowledge, beliefs and suppositions shared by the members of a community. Relative to the personal experiences that lead to physical and linguistic co-presence, cultural co-presence offers less reliable evidence of common ground. Clark and Marshall (1981) argue that the three basic types of definite reference reveal an obvious fit with mutual knowledge: Deixis corresponds with physical co-presence, anaphora with linguistic co-presence, and proper names with cultural co-presence.

In order to efficiently search for evidence of triple co-presence in memory, interlocutors are supposed to keep a mental record of *reference diaries*, including significant events in the personal experiences and cultural histories that they share with others (Clark, 1996; Clark and Marshall, 1978, 1981). These reference diaries are built such that during conversation, interlocutors maintain in memory a model of what their partner has in mind—their perceptions, knowledge and ongoing thoughts, and constantly update it, expanding their common ground with every turn in the conversation. In sum, these reference diaries form the memory basis for physical, linguistic and cultural co-presence, which are the heuristics that interlocutors rely on for establishing mutual knowledge—rather than engaging in an infinite recursion of propositional attitudes.

In a series of laboratory experiments, Clark et al. (1983) tested the hypothesis that, in instances of physical co-presence, listeners understand reference based on what they believe is mutually salient to both interlocutors (see also Wilkes-Gibbs & Clark 1992). However, Keysar (1997) challenged their interpretation of these and other findings, arguing that they are compatible with theories of language comprehension that do not assume a role for common ground. In Keysar's view, an experimental paradigm can only show a reliable effect of common ground during language comprehension when it distinguishes the listener's own perspective from the information they share with the speaker; otherwise, the listener might be using their own, private ground for interpretation, as that would be more parsimonious than tracking mutual knowledge. These arguments mark the birth of the *Director task*, which has become a classic eye-tracking paradigm in psycholinguistics.

In the original setup of the Director task, a participant and a confederate Director stand on opposite sides of a vertical grid of objects, but not all objects are in their common ground because some of the cells are occluded for the Director, who is supposedly ignorant of their contents. Thus, when the Director asks the participant 'Move the small candle,' for example, the smallest of three candles is visible only to the participant, who must then select between the medium-size candle in common ground and the smallest candle in their privileged view. Originally, Keysar et al. (1998a) designed this task to reliably test the claim made by Clark and

Carlson (1981, 328): “The comprehension process must keep track of common ground, and its performance will be optimal if it limits its access to that common ground. Whether its design is actually optimal in this respect is a question that can only be answered empirically.”

Contrary to the *Restricted Search hypothesis* defended by Clark and Carlson (1981), participants in the Director task don't limit their search for the intended referent to those objects in the open cells of the grid. In fact, participants show a tendency to first look at the smallest candle in their privileged view—sometimes even reaching for it, before picking up the medium-sized candle in common ground (Keysar et al., 2000, 1998a,b). Keysar and colleagues interpreted their results as evidence for a two-stage model of reference resolution, whereby listeners' initial search for an intended referent is driven by their own egocentric perspective, and common ground only plays a role in a later monitoring process that detects errors and adjusts interpretation accordingly (Epley et al., 2004). Contrary to Clark and Carlson's view that limiting referent search to common ground would be optimal, Keysar et al. argue that their Egocentric Anchoring and Perspective Adjustment model may be more efficient, since listeners can often afford to be egocentric and rely on their own perception of salience (which should often be analogous to the speaker's) (Keysar et al., 2003) or on the speaker engaging in audience design for their benefit (Keysar et al., 2000), ultimately resorting to conversational repair if need be (Barr and Keysar, 2004).

Pickering and Garrod (2004; 2004) propose another two-stage model of common ground in their interactive alignment theory of dialogue. In their view, Clark's (1996; Clark & Marshall, 1981) psycholinguistic model of common ground is computationally more feasible than the iterated attitudes of mutual knowledge, yet it still requires interlocutors to keep in mind a complex model of their conversational partner's knowledge, separate from their own. Pickering and Garrod argue that, rather than constantly updating a model of their partner's mental states (which is unnecessarily computationally taxing), interlocutors engage in an automatic, unconscious process of interactive alignment, whereby their mental representations are aligned at different levels (e.g., low-level alignment of words and syntax leads to higher-level alignment of situation models). This body of aligned representations, which interlocutors synchronously build as conversation proceeds, is what Pickering and Garrod (2004; 2004) call *implicit common ground*. Interlocutors do not need to infer the contents each other's mind and construct separate mental representations, since they can simply rely on their implicit common ground.

In instances of miscommunication, interlocutors rely on the implicit common ground for interactive repair, which allows them to re-align their representations

without having to model each other's mental states. It is only when these straightforward means of repair fail that interlocutors may go beyond the implicit common ground and rely on *full common ground* (Garrod and Pickering, 2004; Pickering and Garrod, 2004). Unlike implicit common ground, which is the basis of interactive alignment, full common ground is a specialized, non-automatic process that is only used when radical misalignment becomes obvious. Drawing inferences about common ground is therefore an optional strategy that interlocutors employ only when resources allow, and which is normally not required.

In a more recent development of their mechanistic account of dialogue, Pickering and Garrod (2021a; 2021b; Gandolfi et al 2022) introduce the notion of *shared workspace*, which includes all the elements that are in the interlocutors' current focus of attention. For example, in the standard Director task, the objects in the open cells of the grid and the utterances produced by the Director and the participant would all be in their shared workspace. The objects in the occluded cells, on the other hand, would be in the participant's individual space. Rather than being representational, the shared workspace contains the elements over which interlocutors' representations are formed, and alignment is established between their shared representations.<sup>5</sup>

In the shared workspace framework, situation models and linguistic representations are both first-order representations that can form first-order alignments. Yet communicative success also depends on second-order representations, which are invoked when interlocutors meta-represent their alignment. Thus, interlocutors' situation models contain tokens that can be *m-tagged* when interlocutors meta-represent their alignment on a given token (e.g., CANDLE-token-m would correspond to the candle in the middle of the table if both interlocutors were staring at it and/or talking about it). M-tagging also reflects each interlocutor's confidence in their meta-representation of alignment on a continuous scale, such that m-tagged tokens have an m-tag value approaching 1 (e.g., CANDLE-token-m-1 vs LAMP-token-m-0.22). Importantly, m-tagging results from the interlocutors' constant comparison of their respective representations in order to monitor whether they are aligned or not (Gandolfi et al., 2022).

In line with previous psycholinguistic models, Horton and Gerrig (2005a; 2005b; 2016) tried to provide a more psychologically plausible account of common ground than the infinite chain of reasoning proposed by Schiffer (1972) and other for-

---

<sup>5</sup>It is worth noting that psycholinguistic models often refer to the contents of common ground as objects, rather than propositions, although they never seem to make explicit commitments about this view

malists, as well as Clark and Marshall's (1978; 1981) reference diaries, which they deemed too computationally burdensome to form the basis of everyday conversation. In Horton and Gerrig's view, conversational common ground is an emergent property of ordinary processes of memory encoding and retrieval, and does not require specialized mental representations. Thus, each participant in a conversation functions as a memory cue for the retrieval of associated information through an automatic search process known as *resonance*. Rather than maintaining a mental model of the listener, or tagging information to support common ground inferences, interlocutors rely on partner-specific associations in long-term memory (including those between a conversational partner and an intended referent), which work as input for language production processes.

When a conversational partner serves as a memory cue for the retrieval of associated relevant knowledge, utterances are more likely to be consistent with pragmatic expectations about common ground (Horton and Slaten, 2012). However, those memory representations that permit audience design need to become accessible within the time course of language production, and the relative accessibility of particular memory representations depends on people's experiences with their interlocutors. For example, when two people talk about a certain topic regularly, relevant information will be more readily accessible in their next conversation. By contrast, information that is only weakly associated to a conversational partner may not be sufficiently accessible to result in audience design. In those cases where interlocutors do not automatically access information relevant to common ground, they can engage in strategic commonality assessment through more controlled memory search processes. In sum, according to Horton and Gerrig (005a; 005b; 2016), people gain access to the information they share with their interlocutors through domain-general memory processes acting over ordinary memory traces, and these processes may be automatic or strategic depending on the relative accessibility of partner-specific information in long-term memory.

The most recent psycholinguistic theory of common ground was proposed by Heller and Brown-Schmidt (2023), who reject previous notions of common ground in favor of a cognitive architecture that includes separate representations of the interlocutors (self and others), and a cognitive process that continuously compares these representations during conversation, outputting both similarities and differences in perspective. This cognitive architecture is motivated by laboratory tasks where speaker and listener have different visual and/or epistemic perspectives (akin to the Director task), and conversational scenarios with similar constraints, as well as the asking of information questions—a linguistic phenomenon not previously addressed by psycholinguistic models of common ground. Heller

and Brown-Schmidt argue that previous theories characterized communication as the growing of common ground or alignment, losing sight of the importance of differences in the interlocutors' perspectives. In their theory, by contrast, mutual knowledge is just one of the two possible outputs of the comparison process (i.e., perspective similarity), with perspective differences being equally important.

Rather than a dedicated representation for common ground, the Multiple Perspectives Theory posits separate mental state representations for the self and other, a view that has antecedents in the formal tradition (Gunlogson 2004; 2008; Lascarides & Asher 2009; Farkas & Bruce 2010b) and in Pickering and Garrod's shared work-space framework (Gandolfi et al. 2022), which also posits that interlocutors constantly compare their perspectives and metarepresentations of alignment during dialogue. Heller and Brown-Schmidt (2023) conceive of mutual knowledge as a transient representation that updates the interlocutors' representations of self and other and is used in the moment for both language production and comprehension. Based on documented memory limitations and asymmetries between speakers and listeners, Heller and Brown-Schmidt predict that speakers will often assume mutual knowledge when their listener actually needs to rely on inference to derive the intended interpretation.

As the foregoing review shows, psycholinguists have been preoccupied with the tractability of common ground, positing cognitive mechanisms that can account for interlocutors' ability to *quickly* and *efficiently* coordinate their background information. This represents an attempt to lift the main idealization of the formal tradition. However, whereas formalists have sought to expand the empirical scope of their models, psycholinguists have *narrowed* their empirical scope by focusing almost exclusively on definite reference, while mostly ignoring other phenomena that have animated formalists' discussions.

### 3 BRIDGING THE TRADITIONS

The formal and psycholinguistic traditions have become increasingly disconnected—to their detriment, we think. This section aims to remove obstacles to their reintegration.<sup>6</sup> In §3.1, we argue that an apparent substantive disagreement between the two traditions rests, at least in part, on a conceptual misunderstanding. In 3.2, we

---

<sup>6</sup>Some might deny that this is a desirable aim, because they take common ground to be a model of objective and normative features of conversation, rather than features of human psychology (e.g. Gauker 1998; DeVault and Stone 2006; Geurts 2019; Stojnic 2021). We lack the space to properly address these views here, but see Harris (2019; 2024).



will argue that the positive proposals of both traditions are all best understood as compatible, partial solutions to a larger problem.

### 3.1 ITERATED ATTITUDES

Aside from their different idealizations and methodologies, formalists and psycholinguists are divided by an apparently substantive disagreement about iterated-attitude theories of common ground, which tend to be assumed by formalists and rejected by psycholinguists. The origin of this disagreement is Clark and Marshall's (1981, 15) processing paradox, which psycholinguists have tended to cite in agreement (Horton & Gerrig 2005a; Pickering & Garrod 2004; Horton & Brennan 2016). This argument can be seen as inaugurating psycholinguists' search for psychologically realistic alternatives to iterated theories.

We think that this dispute is less substantive than it appears. First, notice that formalists and psycholinguists tend to agree about paradigmatic cases of common ground. For example, Clark's (1996, 93) example of physical co-presence, in which a father and son are looking at a conch shell on the beach, is nearly identical to examples used by Schiffer (1972, 33–36) to illustrate mutual knowledge and by Stalnaker (Stalnaker, 1978, 85) to illustrate how common ground can arise from a *manifest event*.

However, according to Clark and Marshall (1981, 15; cf. 1996), although physical co-presence would justify any of the states of knowledge needed for iterated common knowledge, attaining each of these states would require an extra process of “checking” that the proposition is true. Since this process would take “a finite (though small) amount of time or capacity,” it can't be done an infinite number of times in a finite span. Clark and Marshall therefore propose that interlocutors skip all of those steps by using a heuristic that takes them directly from recognizing physical co-presence to treating information about the conch as common ground.

By contrast, Schiffer and Stalnaker would say that the father and son already possess iterated common knowledge, without needing to go through any extra process. Like many philosophers of mind, they deny that we need to separately check or even entertain each proposition that we come to know. As Schiffer (1972, 36) puts it, “it is no objection to the claim that *S* knows that *p* that the thought that *p* never once entered *S*'s head.” Likewise, they would insist that there is nothing paradoxical about knowing infinitely many propositions: You currently know that your index finger is shorter than five inches long, that it is shorter than six inches long, that it is shorter than seven inches long, and so on, for example, even if you've never gone through any process of checking that each of these propositions is true.

What could knowledge be, such that an infinite amount of it could be acquired by and stored in a finite mind? This, of course, is a matter of controversy (Schwitzgebel, 2023). But, as it happens, the most prominent defenders of iterated theories of common ground are on record defending one of the following three views: (i) *Dispositionalists* argue that to believe (or know) that  $p$  is to be disposed to think and behave in the right ways—e.g., to be disposed to judge  $p$  true if the question arises, or to act on one's plans and desires as if  $p$  is true (Schwitzgebel, 2013). (ii) *Functionalists* argue that to believe that  $p$  is to possess an inner state that causally interacts with perception, action, and other mental states and processes in belief-that- $p$  like ways (Lewis, 1972; Stalnaker, 1984). (iii) *Interpretationists* say that to have a belief is to be someone whose thoughts and behavior can be predicted and explained using a theory that credits one with the belief (Davidson, 1982; Dennett, 1971). By the standards of any of these theories, the father and son in Clark's example each qualify as possessing iterated common belief, and since each of these beliefs is true and epistemically warranted, they also possess iterated common knowledge.

Not every philosopher of mind would agree. (iv) *Representationalists* about belief would argue that to believe (or know) that  $p$  requires having a discrete mental representation of  $p$ —e.g., a suitably tokened sentence in the language of thought (Fodor, 1975; Quilty-Dunn et al., 2022)—and that someone who appears to have infinite beliefs really just has an infinite number of dispositions to quickly acquire the relevant beliefs as needed (Audi, 1994). Turning these dispositions into actual beliefs would require some process that takes time and cognitive capacity. Clark and Marshall's processing paradox therefore seems to tacitly presuppose a representationalist theory of belief/knowledge.

If this diagnosis is correct, then what may have looked like iterated theorists' unresponsiveness to a devastating objection turns out to be a tacit disagreement over the nature of belief that has long divided philosophers of mind. Representationalists take beliefs to be *concrete* mental representations that interact with cognitive mechanisms. By contrast, theories (i)–(iii) take beliefs to be *abstract* properties of minds, which could be realized by a variety of underlying mechanisms. *One* thing that could constitute a belief in the abstract sense is a belief in the concrete sense—e.g., a person's disposition to judge that  $p$  (etc.) could arise from their possession of the right mental representation. But a dispositional belief could also be grounded in a mechanism (such as a heuristic) that can produce a mental representation as needed. This suggests that the mechanisms posited by psycholinguists are compatible with iterated theories. For example, dispositionalists could argue that the psycholinguists' cognitive mechanisms are what give rise to the infinite dispositions that constitute iterated common knowledge. The flip side of this

compatibility, however, is that iterated theorists (i.e., nearly all formalists) should recognize that they have told us little or nothing about the cognitive mechanisms that allow interlocutors to select and represent common ground—a point that has been obscured by their talk of propositional attitudes.

### 3.2 THEORETICAL SCOPE

Our next reconciliatory claim is that no one of the extant theories of common ground can do it all: Each makes sense of at most some of the data that a theory of common ground should explain. Consider again iterated theories. Putting aside questions of psychological plausibility, there are powerful reasons to think that these theories can't account for *every* scenario in which interlocutors distinguish common ground from what they privately believe.<sup>7</sup> One reason is that we regularly communicate in the ways that common ground was posited to explain, even when we know we can't achieve the relevant iterated attitudes with our utterances. This is what Harris calls *publicity-averse* situations (2020; 2025). For example, when you send a message to a colleague who rarely reads their email, or when you teach a class that includes students whom, you think, aren't understanding what you're saying, or when you think that a Parisian waiter didn't understand your order, you don't immediately come to believe (or even accept) that they believe what you said. This means that you and your addressee don't satisfy any of the metarepresentational clauses of an iterated definition of common ground (e.g., anything beyond (1a–b) in (7)).

Crucially, communicators in publicity-averse scenarios can still successfully communicate by means of assertions, and can still use definite reference, anaphora, presupposition, and context-sensitive expressions, in ways that depend on what they previously communicated (Harris, 2020; Simons, 2025). The Parisian waiter may still have understood your order of “steak and french fries,” for example, even if you suspect that they didn't, and in that case, you might still try to communicate more by saying, “Not too much salt on the fries, please.” Of course, this risks miscommunication, but you might be willing to take this risk, given that you have some hope (if not belief) that you will be understood. This suggests that communicators must have some cognitive mechanism or mechanisms for choosing which information to treat as background, other than treating it as commonly known, believed, or accepted.

---

<sup>7</sup>See also Armstrong (2018), who argues that children and some non-human animals, who aren't capable of iterated attitudes, make use of a kind of common ground. This motivates less demanding mechanisms that may also be at work in humans.

These publicity-averse scenarios cause parallel problems for psycholinguistic accounts, as well. For example, Clark's explanation for why the contents of utterances normally become common ground is that they, like the conch shell between the father and son, are physically co-present, and so have the features that trigger the linguistic co-presence heuristic. But publicity-averse situations don't meet this condition, because they give at least one interlocutor salient reasons to doubt that the other can perceive or understand their utterance. These situations also interfere with the *grounding* moves which, Clark (1996, ch.8) thinks, we typically use to signal uptake and reinforce the common ground (e.g., you can't nod in response to an email, and you might suspect that the Parisian waiter's nod does not signal genuine understanding). Similarly, Horton and Gerrig's (2005a; 2016) memory account of common ground does not make the right predictions in those instances where a speaker knows that they shared an experience with their listener (e.g., went to a concert together), but also suspects that the listener might not remember (e.g., because the concert was long time ago). Yet these arguments don't demonstrate that we *never* use iterated attitudes, Clark's heuristics, or Horton and Gerrig's memory-retrieval mechanism as a strategy for selecting which information to treat as common ground: What follows is only that these strategies cannot be successful *in every communicative situation*.

A further threat to the scalability of some accounts is the existence of conversations in which interlocutors treat information they don't believe as common ground. We all routinely participate in conversations of this kind, such as when we feign agreement in order to avoid conflict, or when we deliver a friend to their surprise party. A remarkable feature of these situations is that common ground seems to play its usual role in assertion, presupposition, context-sensitivity, definite reference, and anaphora, irrespective of whether all interlocutors believe everything in the common ground, all are merely pretending to believe, or some believe it and others are pretending. As we explained in §2, formalists have accounted for this by defining common ground in terms of what the interlocutors *accept* for the purpose of the conversation (Stalnaker, 2002, 2014). Acceptance, on this view, is an attitude or category of attitudes that may involve treating a proposition as true *provisionally* and *for some purpose*, such as when one makes a supposition in order to see what follows from it, or in order to play devil's advocate. This is a powerful generalization of the common-ground model, but the mechanisms that psycholinguists have posited for selecting and representing common ground don't seem to be compatible with it.

Clark's physical and linguistic co-presence heuristics depend on what interlocutors actually perceive, and Horton and Gerrig's memory-retrieval mechanism

depends on their actual memories, whereas at least some acceptance-without-belief scenarios would seem to require common ground to depend on our psychological mechanisms for joint pretense or imagination instead. Acceptance-without-belief also solves a problem that Heller and Brown-Schmidt (2023) take to be devastating for the notion of common ground: Namely, that interlocutors can sometimes communicate successfully even when one of them has a false belief (and this is known by the other). However, that argument is only critical if one equates common ground with mutual knowledge, but not if common ground can also include propositions that are merely accepted for the purpose of the conversation. In general, acceptance is a powerful tool for repairing defective contexts.

A further reason to doubt that any extant account can explain all of the data is that information in common ground seems to come in different representational formats. In §2.1, we summarized reasons to think that common ground includes object representations. This line of thought has recently led some researchers to argue that common ground must be at least partially built out of object representations that cognitive scientists have posited for other reasons, such as the *mental files* that philosophers of mind have posited to make sense of object-directed thought, or the *object files* that vision scientists have posited to explain our capacity for object tracking (Brody, 2020; Lewis, 2025; Murez and Recanati, 2016). However, common ground can't be wholly built out of object representations, because some common-ground information isn't about particular objects (e.g., the proposition that all dogs are mammals). This suggests that neither purely propositional *nor* purely objectual representations can do all of the work of common ground. And this line of thought can be extended when we consider other complications of the common-ground model, such as QUDs, which seem to be a representation of neither propositions nor objects, but of *goals* or *plans*. Given its diverse functions, it is an open question how many different kinds of mental representations may be needed to capture common ground.

## 4 COGNITIVE PLURALISM ABOUT COMMON GROUND

We have just argued that no extant theory of common ground can do it all. Nor is it obvious how to generalize any one theory. We think that this is no accident: Humans use a number of different cognitive mechanisms to select and represent common ground. Different researchers have focused on different ones, and have

often understood the results to be competing theories of a single phenomenon. We would like to recast these theories as accounts of mechanisms that might coexist within a single mind, complementing each other as part of a larger, multifaceted solution to the problem of how to select and represent shared background information. We call this view *cognitive pluralism about common ground*, or *pluralism* for short: the idea that we have more than one cognitive mechanism for managing common ground. The advantage of pluralism is that it allows us to account for both the sources of variability in common ground use, and its processing demands. In what follows, we consider each of these advantages in turn.

#### 4.1 EMPIRICAL REASONS FOR PLURALISM: SOURCES OF VARIABILITY

The first source of variability in common ground use stems from the diversity of linguistic phenomena relying on interlocutors' background information. The cognitive mechanisms posited by psycholinguists are customized to explain definite reference, but, as we saw in §§1 and 2.1, common ground also has other jobs to do. It is doubtful that any one cognitive mechanism could take care of the roles of common ground in definite reference, assertion, presupposition (and accommodation), the many forms of context sensitivity, anaphora, and so on—particularly when we consider that information in common ground can be known, believed, or merely accepted for the purpose of the conversation.

Another important source of variability in common ground use comes from the wide range of situations in which people communicate, which makes some psycholinguistic accounts viable in some communicative situations but not in others. For example, Clark and Marshall's (1981) physical co-presence heuristic and Pickering and Garrod's (2012a) shared workspace framework are useful for face-to-face discussions of perceptually available topics, but not when communicators can't see each other or their subject matter. Similarly, Horton and Gerrig's (2005; 2016) automatic memory retrieval mechanism may be useful to draw information from shared memories with interlocutors, but not when talking with strangers, or about topics that aren't connected to shared memories.

A third kind of variability relates to the nature of the mental representations underlying common ground. In §2.1, we showed how work in the formal tradition suggests that common ground needs to encompass information in different representational formats: shared representations of propositions (e.g., the proposition that the incumbent won), representations of the entities under discussion (e.g., the

incumbent), representations of goals (e.g., learning the results of the election), and so on. Since these different kinds of representations have different formats, we need to rely on different cognitive systems in order to make use of them. For example, we may need to make use of object-file systems in order to track individuals under discussion (Brody, 2020; Lewis, 2025), or our planning systems to coordinate goals (Roberts, 2012a).

Relatedly, common ground mechanisms must operate over both short- and long-term memory representations (what Arnold (2008) calls *local vs global assumptions*). However, psycholinguistic models tend to focus on one or the other (cf. Pickering & Garrod 2011a): For example, Horton and Gerrig's (2005a; 2016) model of common ground focuses on long-term memory representations, which could explain the formulation of first mentions (e.g., whether to refer to 'Dr Evans' or 'our children's pediatrician,' depending on whether the addressee knows who Dr Evans is), but could not account for anaphora (e.g., whether to refer to Dr Evans as 'she' depending on its discourse accessibility). Reversely, Heller and Brown-Schmidt's (2023) theory, with its focus on transient representations of mutual knowledge that are used in the moment, could contribute to a theory of how short-term memory figures in anaphora resolution, but not to a theory of how shared memories become accessible in long-term memory for the purpose of formulating first mentions.

Finally, an important source of variability not accounted for by previous theories of common ground are the cross-linguistic differences in grammatical markers of new-familiar information (for discussion, see Shukla et al. 2022; Rubio-Fernandez 2024). For example, whereas speakers of languages with definite articles must signal a familiar referent for the benefit of their addressee (e.g., 'We bought the house' vs 'We bought a house'), in languages without articles, it is often the listener who must infer from the context whether the speaker is referring to something new or familiar. Seen this way, the division of pragmatic labor between speakers and listeners managing common ground may vary depending on the language (Rubio-Fernandez 2025).

#### **4.2 EMPIRICAL REASONS FOR PLURALISM: PROCESSING DEMANDS**

Pluralism is also motivated by considerations of computational complexity or *tractability*, which has played a central role in the psycholinguistic tradition. However, some recent psycholinguistic models raise tractability worries. In Pickering and

Garrod's (2022; 2021a) shared workspace framework, interlocutors need to constantly monitor each other's discourse representations to m-tag those that are aligned, and to do so with an estimated degree of probability. Similarly, in Heller and Brown-Schmidt's (2023) Multiple Perspectives Theory, interlocutors must constantly compare their mental states during conversation, registering all differences and similarities, also in probabilistic terms. It seems to us right that interlocutors must at least sometimes represent discrepancies in their background information—e.g., to decide whom to ask a question, identify and repair defective contexts, or accommodate presuppositions (Simons, 2025).<sup>8</sup> However, there are good reasons to think that representing conflicting beliefs is cognitively demanding (Westra and Nagel, 2021), and it is therefore questionable whether interlocutors can communicate while constantly comparing their mental states and classifying them as similar or different with a distinct probability—especially in multi-party conversation.<sup>9</sup>

Another important consideration regarding the processing demands of using common ground is the relative efficiency of different psycholinguistic accounts. Keysar et al.'s (2000; 1998a; 2003) Egocentric Anchoring and Adjustment model and Pickering and Garrod's (2004) original interactive alignment mechanism were put forward as efficient alternatives to the cognitive demands of constantly monitoring common ground in conversation. While both these models allowed for a late, effortful use of common ground for conversational repair, they do not endorse cognitive pluralism, in our sense of the term, since they don't consider egocentric processing as an efficient mechanism for managing common ground (but rather as an *alternative* to using common ground).

Arnold (2008; 2010) also put forward an efficient referent production system whereby speakers use their own discourse model to calculate the relative accessi-

---

<sup>8</sup>Formal-pragmatic models typically idealize away from phenomena that involve representing differences in perspective. However, formalists were the first to acknowledge that these are real and important phenomena that a theory of common ground must explain, identifying defective contexts (Stalnaker, 1978) and publicity-averse situations (Harris, 2020). If formalists haven't modeled perspective differences yet, it is only because nobody currently knows how to create predictive formal models of these phenomena.

<sup>9</sup>Heller and Brown-Schmidt propose that this continuous comparison process is limited to *relevant* information, but they don't operationalize relevance and cite several authors whose concepts of relevance are neither compatible with each other, nor with their own view: The relevance theoretic notion is the ratio between cognitive effects and processing effort (hence a technical term for *efficiency*) (Sperber and Wilson, 1995), while Grice's and Roberts' notion of relevance is indeed the one intended by Heller and Brown-Schmidt (meaning *pertinence*). Yet Grice (1975) offers no formalisation of this notion, and the one proposed by Roberts (2004) relies on the very notion of common ground that Heller and Brown-Schmidt are trying to eschew.



bility of a given referent (e.g., whether Dr Evans was mentioned recently enough to refer to her as ‘she’), rather than explicitly considering the addressee’s discourse model. Along similar lines to Keysar et al. and Pickering and Garrod, Arnold argues that this efficient mechanism would often have the same outcome as taking the addressee’s perspective, while being less cognitively demanding. However, Arnold acknowledges that her prediction crucially depends on the addressee being attentive during the conversation (otherwise they might not understand who ‘she’ refers to).

From a pluralist perspective, a continuous comparison of the interlocutors’ mental representations (Gandolfi et al., 2022; Heller and Brown-Schmidt, 2023; Pickering and Garrod, 2021a) would be an extremely taxing default mechanism, whereas an egocentric language production/ comprehension system (Arnold, 2008, 2010; Garrod and Pickering, 2004; Keysar et al., 2000, 1998a, 2003; Pickering and Garrod, 2004) would often fail depending on the listener’s engagement in the conversation (which can compensate for speakers’ egocentric language production) and the speaker’s audience design (which can compensate for listeners’ egocentric language comprehension). Instead of proposing default strategies operating across all communicative situations, a cognitive-pluralist account would posit a variety of common ground mechanisms: those that project our own perspective as shared with our interlocutor—and are therefore more efficient, and those that require actively taking our interlocutor’s perspective—which are more costly, but often worthwhile.

Finally, cognitive pluralism is compatible with the view that language users not only have a variety of cognitive mechanisms at their disposal to manage common ground in different communicative situations and for different linguistic phenomena, but they can actually *combine* different cognitive mechanisms in any one exchange depending on their goals, what is at stake, and their cognitive resources. For example, if two Canadians see a hockey puck sitting between them, physical co-presence may be enough to make it common ground that *something* is there, but in order to make it common ground *that it is a hockey puck*, they may also need to rely on cultural co-presence. Yet these strategies can also be profitably combined with explicit meta-representation. Thus, one of our two Canadians, having deployed both heuristics, might also wonder whether the other doubts their knowledge of hockey (since they themselves don’t look stereotypically Canadian). This sort of meta-representation could be an optional way to second-guess or fine-tune a representation of common ground that has been built using rough-and-ready heuristics. Of course, as we saw in §3.2, publicity-averse situations (e.g., when we email a colleague who rarely reads their email) prevent double checking common

ground. Therefore, meta-representation, too, is best seen as just one more cognitive mechanism, which is useful in some situations and for some purposes, but not others.

## 5 CONCLUSIONS AND FUTURE DIRECTIONS

Common ground emerged as part of a model that offered unifying explanations of diverse pragmatic phenomena. The formal and psycholinguistic traditions have each lifted some of this model's idealizations, and have thereby made significant progress, but in orthogonal directions: The formal tradition has built more complicated formal models to explain more phenomena, while ignoring questions about cognitive mechanisms. The psycholinguistic tradition theorized increasingly realistic cognitive mechanisms while narrowing the range of phenomena they explained.

We have tried to begin the process of integrating these two orthogonal forms of progress. We hope that future theories of common ground will incorporate *both* broad empirical coverage *and* realistic, independently motivated theories of the underlying psychological mechanisms. We think that these theories would do well to self-consciously adopt a cognitive-pluralist approach, which would in turn lead them to seek new empirical methodologies that can uncover the different mixtures of cognitive mechanisms that support our capacity to select and represent common ground in different situations.

## References

- Armstrong, J. (2018). Provincialism in pragmatics. *Philosophical Perspectives*, 32(5–40).
- Arnold, J. E. (2008). Reference production: Production-internal and addressee-oriented processes. *Language and cognitive processes*, 23(4):495–527.
- Arnold, J. E. (2010). How speakers refer: The role of accessibility. *Language and Linguistics Compass*, 4(4):187–203.
- Asher, N. and Lascarides, A. (2001). Indirect speech acts. *Synthese*, 128(1/2):183–228.

- Audi, R. (1994). Dispositional beliefs and dispositions to believe. *Nous*, 28(4):419–434.
- Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, 25:1–36.
- Barr, D. J. and Keysar, B. (2004). Making sense of how we make sense: The paradox of egocentrism in language use. In *Figurative language comprehension*, pages 21–41. Routledge.
- Beaver, D., Roberts, C., Simons, M., and Tonhauser, J. (2017). Questions under discussion: Where information structure meets projective content. *Annual Review of Linguistics*, 3:265–284.
- Bledin, J. and Rawlins, K. (2018). Resistance and resolution in discourse. *Unpublished Manuscript*.
- Brody, G. (2020). *Indexing Objects in Vision and Communication*. PhD thesis, Central European University, Budapest.
- Camp, E. (2018). Insinuation, common ground, and the conversational record. In Fogal, D., Harris, D., and Moss, M., editors, *New Work on Speech Acts*, pages 40–66. Oxford University Press, Oxford.
- Carlson, L. (1982). *Dialogue Games: An Approach to Discourse Analysis*. Number 17 in Synthese Language Library. D. Reidel, Dordrecht, The Netherlands.
- Clark, H. H. (1996). *Using Language*. Cambridge University Press, Cambridge.
- Clark, H. H. and Carlson, T. B. (1981). Context for comprehension. In Long, J. and Baddeley, A., editors, *Attention and Performance*, chapter 18, pages 313–330. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Clark, H. H. and Marshall, C. R. (1978). Reference diaries. In Waltz, D. L., editor, *Theoretical issues in natural language processing*, volume 2, pages 57–63. New York.
- Clark, H. H. and Marshall, C. R. (1981). Definite reference and mutual knowledge. In Joshi, A. K., Webber, B. L., and Sag, I. A., editors, *Elements of Discourse Understanding*, pages 10–63. Cambridge University Press, Cambridge.

- Clark, H. H., Schreuder, R., and Buttrick, S. (1983). Common ground and the understanding of demonstrative reference. *Journal of Verbal Learning and Verbal Behavior*, 22:245–258.
- Davidson, D. (1982). Rational animals. *Dialectica*, 36:317–327.
- Dennett, D. C. (1971). Intentional systems. *The Journal of Philosophy*, 68(4):87–106.
- DeVault, D. and Stone, M. (2006). Scorekeeping in an uncertain language game. In Schlangen, D. and Fernández, R., editors, *Brandial 2006: Proceedings of the 10th Workshop on the Semantics and Pragmatics of Dialogue*, volume 10, pages 139–146. SemDial, University of Potsdam.
- Dever, J. and Schiller, H. I. (2021). This paper might change your mind. *Noûs*, 55(4):863–890.
- Epley, N., Keysar, B., Van Boven, L., and Gilovich, T. (2004). Perspective taking as egocentric anchoring and adjustment. *Journal of personality and social psychology*, 87(3):327.
- Farkas, D. F. and Bruce, K. B. (2010a). On reacting to assertions and polar questions. *Journal of Semantics*, 27:81–118.
- Farkas, D. F. and Bruce, K. B. (2010b). On reacting to assertions and polar questions. *Journal of semantics*, 27(1):81–118.
- Fodor, J. A. (1975). *The Language of Thought*. Thomas Y. Crowell Company, New York.
- Gandolfi, G., Pickering, M. J., and Garrod, S. (2022). Mechanisms of alignment: shared control, social cognition and metacognition. *Philosophical Transactions of the Royal Society B*, 378(1870):20210362.
- Garrod, S. and Pickering, M. J. (2004). Why is conversation so easy? *Trends in cognitive sciences*, 8(1):8–11.
- Gauker, C. (1998). What is a context of utterance. *Philosophical Studies*, 91:149–172.
- Geurts, B. (2019). Communication as commitment sharing. *Theoretical Linguistics*, 45(1–2):1–30.

- Gibbard, A. (2003). *Thinking How to Live*. Harvard University Press.
- Grice, H. P. (1975). Logic and conversation. In Cole, P., editor, *Syntax and Semantics 3: Pragmatics*. Academic Press, New York.
- Groenendijk, J. and Stokhof, M. (1991). Dynamic predicate logic. *Linguistics and Philosophy*, 14:39–100.
- Gunlogson, C. (2004). *True to form: Rising and falling declaratives as questions in English*. Routledge.
- Gunlogson, C. (2008). A question of commitment. *Belgian Journal of Linguistics*, 22(1):101–136.
- Hamblin, C. L. (1973). Questions in montague english. *Foundations of Language: International Journal*, 10(1):41–53.
- Harris, D. W. (2019). Intention and commitment in speech acts. *Theoretical Linguistics*, 45(1–2):53–67.
- Harris, D. W. (2020). We talk to people, not contexts. *Philosophical Studies*, 177(9):2713–2733.
- Harris, D. W. (2024). review of *Context and Coherence: The Logic and Grammar of Prominence* by una stojnić. *Philosophical Review*, 133(1):87–91.
- Harris, D. W., Fogal, D., and Moss, M. (2018b). Speech acts: The contemporary theoretical landscape. In Fogal, D., Harris, D. W., and Moss, M., editors, *New Work on Speech Acts*. Oxford University Press, Oxford.
- Heal, J. (1978). Common knowledge. *The Philosophical Quarterly*, 28(111):116–131.
- Heim, I. (1982). *The Semantics of Definite and Indefinite Noun Phrases*. PhD Dissertation, University of Massachusetts at Amherst.
- Heim, I. (1983a). File change semantics and the familiarity theory of definiteness. In Bäuerle, R., Schwarze, C., and von Stechow, A., editors, *Meaning, Use and Interpretation of Language*, pages 164–189. de Gruyter, Berlin.
- Heim, I. (1983b). On the projection problem for presuppositions. In *Second West Coast Conference on Formal Linguistics (WCCFL 2)*, pages 114–125.

- Heller, D. and Brown-Schmidt, S. (2023). The multiple perspectives theory of mental states in communication. *Cognitive Science*, 47(7):e13322.
- Hoek, D. (2018). Conversational exculpation. *The Philosophical Review*, 127(2):151–196.
- Horton, W. S. and Brennan, S. E. (2016). The role of metarepresentation in the production and resolution of referring expressions. *Frontiers in Psychology*, 7:1–12.
- Horton, W. S. and Gerrig, R. J. (2005a). Conversational common ground and memory processes in language production. *Discourse Processes*, 40(1):1–35.
- Horton, W. S. and Gerrig, R. J. (2005b). The impact of memory demands on audience design during language production. *Cognition*, 96(2):127–142.
- Horton, W. S. and Gerrig, R. J. (2016). Revisiting the memory-based processing approach to common ground. *Topics in Cognitive Science*, 8(4):780–795.
- Horton, W. S. and Slaten, D. G. (2012). Anticipating who will say what: The influence of speaker-specific memory associations on reference resolution. *Memory & cognition*, 40:113–126.
- Karttunen, L. (1974). Presupposition and linguistic context. *Theoretical Linguistics*, 1:181–194.
- Karttunen, L. (1976). Discourse referents. In McCawley, J. D., editor, *Syntax and Semantics 7: Notes from the Linguistic Underground*, pages 363–385. Academic Press.
- Kaufmann, M. (2012). *Interpreting Imperatives*. Springer, Dordrecht.
- Keysar, B. (1997). Unconfounding common ground. *Discourse Processes*, 24(2–3):253–270.
- Keysar, B., Barr, D. J., Balin, J. A., and Brauner, J. S. (2000). Taking perspective in conversation: The role of mutual knowledge in comprehension. *Psychological Science*, 11(1):32–38.
- Keysar, B., Barr, D. J., Balin, J. A., and Paek, T. S. (1998a). Definite reference and mutual knowledge: Process models of common ground in comprehension. *Journal of Memory and Language*, 39(1):1–20.

- Keysar, B., Barr, D. J., and Horton, W. S. (1998b). The egocentric basis of language use: Insights from a processing approach. *Current Directions in Psychological Science*, 7(2):46–50.
- Keysar, B., Lin, S., and Barr, D. J. (2003). Limits on theory of mind use in adults. *Cognition*, 89:25–41.
- Langton, R. (2012). Beyond belief: Pragmatics in hate speech and pornography. In Maitra, I. and McGowan, M. K., editors, *Speech and Harm: Controversies over Free Speech*, pages 72–93. Oxford University Press, Oxford.
- Lascarides, A. and Asher, N. (2009). Agreement, disputes and commitments in dialogue. *Journal of semantics*, 26(2):109–158.
- Lederman, H. (2017). *Common Knowledge*, chapter 14, pages 181–195. Routledge, London and New York.
- Lewis, D. K. (1969). *Convention: A Philosophical Study*. Harvard University Press.
- Lewis, D. K. (1972). Psychophysical and theoretical identifications. *Australasian Journal of Philosophy*, 50(3):249–258.
- Lewis, D. K. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(3):339–359.
- Lewis, K. (2014). Do we need dynamic semantics? In Burgess, A. and Sherman, B., editors, *Metasemantics*. Oxford University Press.
- Lewis, K. (2021). Anaphora and negation. *Philosophical Studies*, 178(5):1403–1440.
- Lewis, K. (2025). Discourse referents in a dynamic pragmatics. In Lepore, E. and Sosa, D., editors, *Oxford Studies in Philosophy of Language*, volume 4. Oxford University Press, Oxford and New York.
- MacFarlane, J. (2016). Vagueness as indecision. *Proceedings of the Aristotelian Society, Supplementary Volume*, 90(1):255–283.
- Mandelkern, M. (2024). *Bounded Meaning: The Dynamics of Interpretation*. Oxford University Press, Oxford.
- McGowan, M. K. (2019). *Just Words: On Speech and Hidden Harm*. Oxford University Press, Oxford and New York.

- Murez, M. and Recanati, F. (2016). Mental files: An introduction. *Review of Philosophy and Psychology*, 7(2):265–281.
- Pickering, M. and Garrod, S. (2021b). The shared-workspace framework for dialogue and other cooperative. *Human-like machine intelligence*, page 260.
- Pickering, M. J. and Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and brain sciences*, 27(2):169–190.
- Pickering, M. J. and Garrod, S. (2021a). *Understanding dialogue: Language use and social interaction*. Cambridge University Press.
- Portner, P. (2007). Imperatives and modals. *Natural Language Semantics*, 15:351–383.
- Quilty-Dunn, J., Porot, N., and Mandelbaum, E. (2022). The best game in town: The re-emergence of the language of thought hypothesis across the cognitive sciences. *Behavioral and Brain Sciences*, forthcoming.
- Roberts, C. (2004). Context in dynamic interpretation. In Horn, L. and Ward, G., editors, *The Handbook of Pragmatics*, pages 197–220. Blackwell.
- Roberts, C. (2012a). Information structure: Afterward. *Semantics and Pragmatics*, 5(7):1–19.
- Roberts, C. (2012b). Information structure in discourse: Toward an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5:1–69.
- Rubio-Fernandez, P. (2024). Cultural evolutionary pragmatics: Investigating the codevelopment and coevolution of language and social cognition. *Psychological Review*, 131(1):18.
- Rubio-Fernandez, P. (2025). First acquiring articles in a second language: A new approach to the study of language and social cognition. *Lingua*, 313:103851.
- Schaffer, J. (2007). Knowing the answer. *Philosophy and Phenomenological Research*, 75(2):383–403.
- Schelling, T. C. (1960). *The Strategy of Conflict*. Harvard University Press, Cambridge, MA.
- Schiffer, S. (1972). *Meaning*. Oxford University Press, Oxford.



- Schoubye, A. and Stokke, A. (2016). What is said? *Noûs*, Online First.
- Schwitzgebel, E. (2013). A dispositional approach to the attitudes: Thinking outside of the belief box. In Nottelman, N., editor, *New Essays on Belief*, chapter 75-99. Palgrave Macmillan, London.
- Schwitzgebel, E. (2023). Belief. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Stanford University, <<https://plato.stanford.edu/entries/belief/>>.
- Shukla, V., Long, M., and Paula, R.-F. (2022). Children's acquisition of new/given markers in english, hindi, mandinka and spanish: Exploring the effect of optionality during grammaticalization. *Glossa Psycholinguistics*, 1(1).
- Simons, M. (2001). Disjunction and alternativeness. *Linguistics and Philosophy*, 24:597-619.
- Simons, M. (2025). Availability without common ground. *Linguistics and Philosophy*.
- Simons, M., Beaver, D., Roberts, C., and Tonhauser, J. (2017). The best question: Explaining the projection behavior of factives. *Discourse Processes*, 54(3):187-206.
- Sperber, D. and Wilson, D. (1995). *Relevance: Communication and Cognition*. Blackwell, Oxford.
- Stalnaker, R. (1970). Pragmatics. *Synthese*, 22:272-289.
- Stalnaker, R. (1973). Presuppositions. *Journal of Philosophical Logic*, 2:447-457.
- Stalnaker, R. (1974). Pragmatic presuppositions. In Munitz, M. and Unger, P., editors, *Semantics and Philosophy*, pages 197-214. New York University Press, New York.
- Stalnaker, R. (1978). Assertion. In Cole, P., editor, *Syntax and Semantics 9*, pages 315-332. Academic Press, New York.
- Stalnaker, R. (1984). *Inquiry*. MIT Press, Cambridge, Mass.
- Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25:701-721.
- Stalnaker, R. (2014). *Context*. Context and Content. Oxford University Press, Oxford.

- Stojnić, U. (2021). *Context and Coherence: The Logic and Grammar of Prominence*. Oxford University Press, Oxford.
- Stokke, A. (2018). *Lying and Incincerity*. Oxford University Press, Oxford.
- Stokke, A. (2023). Fictional force. *Philosophical Studies*, 180:3099–3120.
- Westra, E. and Nagel, J. (2021). Mindreading in conversation. *Cognition*, 210(104618):1–15.
- Wilkes-Gibbs, D. and Clark, H. H. (1992). Coordinating beliefs in conversation. *Journal of memory and language*, 31(2):183–194.
- Yalcin, S. (2007). Epistemic modals. *Mind*, 116(464):983–1026.
- Yalcin, S. (2012). Bayesian expressivism. *Proceedings of the Aristotelian Society*, 112(2):123–160.
- Yalcin, S. (2018). Belief as question-sensitive. *Philosophy and Phenomenological Research*, 97(1):23–47.
- Yalcin, S. (2024). Defining common ground. *Linguistics and Philosophy*, 47:1045–1070.