Genre and Conversation

Elmar Unnsteinsson and Daniel W. Harris

Draft of September 20, 2024

Abstract

Conversations can belong to different genres. Some are about sharing information while others are about making decisions; some are about making firm commitments while others are about brainstorming options; some are about sticking to the facts while others involve make-believe; some are highly cooperative while others are adversarial. These are orthogonal dimensions of variation, each of which makes some kinds of speech acts more natural and expected than others. So then, what are these genres, and why do they exist? We argue that these four dimensions of genre variation can be understood in terms of commonly recurring patterns in conversation plans, which are the structures of intentions that we use to organize conversations. Within conversation plans, we find shared intentions to address questions, but also intentions about how to address them, and shared intentions about what kind of information we're presupposing. Conversation plans are also shaped by interlocutors' unshared intentions for the conversation, which may be more or less compatible with those of other participants. We argue that different dimensions of genre variation must be understood in terms of these different kinds of elements in conversation plans. We show that some standard pragmatic tools, including Grice's maxims, are genre-specific. We then argue that genres are useful

*Authors made equal contributions and are listed in random order. For feedback on earlier drafts, thanks to audiences at the CUNY Graduate Center, the Iceland Meaning Workshop, Northwestern University, Jagiellonian University, University of Iceland, and University College Dublin. In particular, thanks to Sam Berstler, Justin Bledin, Ray Buchanan, Elisabeth Camp, Peter van Elswyk, Sanford Goldberg, Steve Hernandez, Leopold Hess, Izabela Skoczeń, Karen Lewis, Eleonora Maccarone, Cornelia Mayer, Devin Morse, Eleonore Neufeld, Ethan Nowak, Katherine O'Donnell, Griffin Pion, Krzysztof Poslajko, Craige Roberts, Paula Rubio-Fernandez, Stephen Schiffer, and Henry Schiller. This work was supported by the Icelandic Research Fund (grant no. 2410058).

because conversation plans provide us with significant sources of nonlinguistic evidence about speakers' intentions, allowing us to communicate more complex messages with less work.

1 Introduction

Suppose that Sandra says to Abe, "I love Italian cinema." What is Sandra's point? Without more information, it's hard to tell. One thing that would be useful would be to know what kind of conversation they are having. If they are merely comparing their movie preferences, Sandra's utterance can be taken as an assertion, but if they are having a practical conversation about what to do tonight, it would be natural to interpret it as an indirect suggestion or request. Likewise, it would be nice to know whether, at this stage of the conversation, Sandra and Abe are merely floating options about what to do or trying to narrow down the options they've already considered. Are they discussing their actual preferences, for that matter, or engaging in some form of collective pretense? Finally, are they having an antagonistic argument, in which case Sandra's utterance might count as a defensive rebuke, or is their conversation a congenial and cooperative one, in which case it is a friendly observation or suggestion?

Each of these questions pertains to what we'll call "conversational genre." A conversation's genre makes certain speech acts more natural than others, and shapes expectations about what others are likely to do with their words. But what are genres, and why do they exist?

We will propose to think of genres as recurring patterns in conversation plans, which are the structures of intentions that govern our communicative exchanges. For example, a practical, decision-making conversation is one that is governed by the interlocutors' shared plan to make certain decisions, whereas an informational conversation is governed by the interlocutors' overarching shared plan to share information. Certain kinds of conversation plans make certain kinds of speech acts most natural because we form communicative intentions as a way of furthering our broader plans, and we are under rational pressure to do so in a coherent and constructive way. Because a conversation's genre licenses interlocutors' expectations about the kinds of intentions their interlocutors can be expected to have, genres facilitate efficient communication.

2 Conversational genre: a first pass

Conversational genres have mostly been discussed within sociolinguistics. For example, Dell Hymes tells us that "genres…are categories such as poem, myth, tale, proverb, riddle, curse, prayer, oration, lecture, commercial, form letter, editorial, etc" (1977, 61). Sociolinguists have often focused on particular, local examples, such as sports-announcer talk (Ferguson 1983), wedding invitations in Brunei (Clynes and Henry 2004), and doctor-patient interviews (Coulthard and Ashby 1975). These discussions tend not to include much theory about the nature of genre, where it comes from, why it exists, and how it shapes conversations. And they tend to focus on conventionalized genres, involving cultural scripts with specified roles and rituals.[1]

We think that these conventional genres are real and interesting. But we will focus on four dimensions of genre difference that are not inherently conventional:

(1) **Informative vs. practical conversations**
Are we sharing information or making decisions?

(2) **Committal vs. exploratory conversations**
Are we taking on commitments or merely floating options?

(3) **Factual vs. make-believe conversations**
Is this how things actually are or are we making it up?

(4) **Cooperative vs. adversarial conversations**
Are our goals complimentary or in conflict?

By describing these genres as non-conventional, we don't mean to deny that each speech community has its own conventions for navigating them—e.g., for signaling that they are engaging in a practical conversation as opposed to an informative one. But these are conventions of implementation rather than of the nature of the genres themselves. We assume that all speech communities engage in conversations that vary along each of these dimensions, whereas sports-announcer talk and doctor-patient interviews owe their very existence to culturally local conventions.

[1]Biber and Conrad (2019, 23–4) are an exception, since they consider the extent to which variation in genre and register might be a linguistic universal.

Our four genre axes are orthogonal. Consider just one permutation, in which two people playing Dungeons and Dragons brainstorm about what to do next within the game: "We could fight the orc." This conversation belongs to the *make-believe*, *exploratory*, *practical*, and *cooperative* sides of these axes. However, the conversation might later become more committal if they begin making decisions rather than exploring options, more informative if they begin merely discussing the properties of orcs, or more adversarial if they begin disagreeing about what to do. And we can imagine a similar range of variation arising in a factual conversation about a real-life aggressor, rather than an orc.

Notice, also, that the genre may be more or less undecided or unclear. Suppose that you run into an acquaintance on the train and begin to talk. Will this be an informative conversation about what's happening in your lives, or a practical one about when to meet for lunch? If the latter, will you commit to a day to meet, or merely float some options? Will you have a serious, factual conversation, or just joke around? Will this be a friendly conversation, or might it have an edge of conflict to it? In principle, each of these questions could be unsettled, in which case the conversation doesn't (yet) belong to any of the genre categories we've outlined. We take it as a datum that such conversations exist, that they admit of a greater degree of conversational freedom, or open-endedness, than genre-bound conversations, and also that this open-endedness can sometimes make such conversations more awkward to navigate, as interlocutors have fewer expectations about what will come next.

As these examples illustrate, conversations can switch back and forth between different genres at different times. We will focus, therefore, on the genre of a conversation *at a time*. At one moment, a conversation may be cooperative and exploratory, only to become contentious and committal later on. Ultimately, we would like to understand how speakers navigate and effect such changes, although we have only limited things to say about this here.

We have said that a conversation's genre makes certain speech acts more natural and others less so. We will illustrate this with respect to each of our dimensions in the subsections below. But first, a disclaimer: We don't mean to suggest that genres make any speech acts impossible. For example, it is possible to say merely informative things in practical conversations. However, we think that genre-defying speech acts typically give rise to a greater risk of miscommunication, because they contravene interpreters' expectations, and so these speech acts typically require more ingenuity on the part of the speaker in order to override these expectations, and may sometimes have to be marked for this purpose. For example, in a practical conversation, a speaker might accompany a merely

informative remark by saying, "I'm not sure what the practical upshot of that is."

Philosophers and linguists have tended to work with models that treat informative, committal, factual, and cooperative conversations as paradigmatic, effectively focusing on the items on the left of our list and idealizing away from the items to the right.[2] For example, we'll argue in §8 that although Grice's cooperative principle is genre-neutral, some of his maxims of conversation are genre specific, and the same can be said of Williamson's (2000) knowledge norm of assertion.

In the remainder of this section, we will zoom in on the four genre distinctions mentioned above, in order to get a pre-theoretic grip on them.

2.1 Informative vs. practical

In informative conversations, we share information, whereas in practical conversations, we make decisions. In practical conversations, the default speech acts are directives, suggestions, proposals, and questions about what to do. Assertions are natural in a practical conversation only insofar as they bear on the decision at hand, and tend to be interpreted as indirect directives or practical suggestions. For example, suppose that we are deciding where to have dinner and I utter (5):

(5) Calexico has good tacos.

Although (5) is a declarative sentence, whose canonical function is to perform an assertion, you will probably interpret my utterance as an indirect proposal to eat at Calexico.

Informative conversations can have a similar influence on how our speech acts are interpreted. Consider the following exchange:

(6) (a) Anya: I am not trying to get to Grand Central Station, but I am curious: how does one get there?
(b) Baker: Take the 5 train.

[2] There are some noteworthy exceptions. One is Yalcin's (2007) discussion of "conversational tone," which we will incorporate into our own view in §6. Another is Clark's (1996) work on dimensions of variation in "conversational setting," although he is interested in kinds of variation other than those on which we focus here. And philosophers have long been interested in various features of non-cooperative speech—slurs, hate speech, silencing—although it is only very recently that some have focused on adversarial interactions as such (see, e.g., Cappelen and Dever 2019, McGowan 2019, McKinney 2016, Camp 2018).

Here Baker utters an imperative sentence—one that would normally be used to perform either a directive act, such as a command, or a practical suggestion. But in this context, Baker manages to merely offer some information, apparently because Anya has made it clear that this is an informative conversation. This is what some authors have called an "instructional" use of an imperative—one whose point is informative rather than directive (e.g. Kaufmann 2012, 141). This is an atypical use of an imperative, which some theories of imperatives have struggled to make sense of (Condoravdi and Lauer, 2012).

Finally, consider sub-clausal utterances, such as (7) whose illocutionary force is not indicated by any clause-level features.

(7) Calexico.

If uttered in an informative conversation—e.g., one where we are trying to figure out which restaurants serve fish tacos—an utterance of (7) is most naturally understood as an assertion. But if we're having a practical conversation—say, one about where to eat dinner—the same utterance is most naturally understood as a practical proposal, elliptical for something like, "let's eat at Calexico."

2.2 Committal vs. exploratory

In committal conversations, utterances are treated as firm commitments. One can respond to an assertion by asking the speaker how they know, or criticize them for saying something false, or unlikely, or for which they lack evidence (Carter, 2024; Grice, 1975; Williamson, 1996, 2000). By contrast, in exploratory conversations, such as brainstorming sessions, we merely explore possibilities without committing ourselves to them. A useful characterization of normative exploratory conversation comes from the influential negotiation textbook, *Getting to Yes* (Fisher and Ury, 1981, 60):

> A brainstorming session is designed to produce as many ideas as possible to solve the problem at hand. The key ground rule is to postpone all criticism and evaluation of ideas. The group simply invents ideas without pausing to consider whether they are good or bad, realistic or unrealistic. With those inhibitions removed, one idea should stimulate another, like firecrackers setting off one another.

As an example, consider a group of doctors brainstorming about a difficult diagnosis, writing their proposals on a whiteboard—a conversation of a kind that took place in most episodes of the TV series, *House*:

(8) Thirteen: A brain tumor is restricting blood flow to her hypothalamus.
Kutner: She fell and hit her head.
Taub: She has lupus.
Thirteen: A toxin.

We can assume that the doctors take the patient's symptoms to have, at most, one of these causes. Still, in this exploratory conversation, it is permissible to float incompatible possibilities. The doctors' utterances almost seem to have an implicitly modal flavor, as if they were merely claims about what *might* be causing the patient's symptoms, as opposed to assertions about what the causes in fact are. But in this unambiguously exploratory conversation, they needn't actually include possibility modals in order to take on this just-floating-a-possibility flavor. Once again, this is also true of sub-clausal utterances, such as Thirteen's. By contrast, in committal conversations, of the kind that most pragmatic theories attempt to model, it is infelicitous to utter declarative sentences with inconsistent contents.

Of course, it's not as if *any* possibility should be floated in an exploratory conversation. If one of our doctors suggested that the patient is suffering from a demonic curse, this would likely be treated by the others as beyond the pale. There are usually limits to what we are willing to treat as live options.

2.3 Factual vs. make-believe

In factual conversations, we try to say things that we regard as true of the actual world, and we may hold others responsible if they fail to live up to this ideal. But when we engage in make-believe—telling fictional stories or engaging in pretend play, for example—the relevant norms of truth and knowledge are not in place. Imagine a parent who says to their child, "There was once a woman who had a third eye in her bellybutton; her name was Gilitrutt." It would not fit the genre of this conversation to reply, "How do you know?" or "That's false." According to one family of theories of fictional discourse, this is because the parent's aim is for the child to join them in taking the attitude of pretense, or make-believe, rather than belief, to the proposition being expressed.[3]

[3] We are assuming a theory of fictional discourse on which fictional and factual speech acts may share content but always differ with respect to the attitudes that speakers adopt, and intend their addressees to adopt, toward those contents. For versions of this sort of attitude-first theory, see Currie (1986, 1990); Searle (1975); Stokke (2023). Others have argued that fictional and factual discourse both involve interlocutors taking the same attitudes (e.g., ordinary belief) toward

Make-believe conversations are united by the fact that interlocutors presuppose and assert information that they do not regard as factual. Aside from fictional storytelling and pretend play, there are other interesting examples of conversations of this kind. Consider, for example, a conversation in which two people take on some supposition that they don't believe in order to explore its consequences. For example, imagine that Alvy and Allison are having a conversation about John F. Kennedy's assassination, and they have temporarily taken on the supposition that a second shooter, in addition to Lee Harvey Oswald, was involved. Allison might say, "Everyone was in on it: Earl Warren, Lyndon Johnson, the CIA." Here, Allison is not attempting to assert anything factual, but is putting forward a proposition that would be true if the supposition were true (holding as much else fixed as possible).

2.4 Cooperative vs. adversarial

There are probably a number of different senses in which conversations can be more or less cooperative. Here we will focus on the sense in which, in cooperative conversations, we can count on our interlocutors to be forthcoming with relevant information, to interpret our utterances charitably, and to make judicious use of our time and attention. Adversarial conversations, such as a tense deposition or a fight with a spouse, aren't like this. In adversarial conversations, it is risky to use ambiguous expressions or to communicate indirectly, as an addressee will be less likely to attribute the intended meaning and more likely to attribute an unintended one.

In some adversarial contexts—typically, those in which everyone knows about the sources and nature of the conflict—we are less likely to treat the same utterance as carrying an implicature. For example, consider the following exchange:

(9) A: Were you at the Tipsy Elf at noon yesterday?
B: No.

In a cooperative context—say, a conversation in which friends are casually chatting about what they had done on the previous day—it would be natural for A to interpret B as implicating that they weren't at the Tipsy Elf at any time close to noon yesterday. After all, if B had been there at half past noon, and they

different contents (Lewis, 1978; Thomasson, 1998). We will assume an attitude-based theory in the main text, here and below, but see footnote 16 for a brief explanation of how our theory of conversational genre could be adapted to account for content-based theories.

think that this might be relevant to A, they probably would have volunteered this information—perhaps as follows:

(10) A: Were you at the Tipsy Elf at noon yesterday?
B: No. But I was there at half past noon.

In a cooperative conversation, not to volunteer this information suggests that B intended A to infer that nothing of the kind was true. If A later learned that B arrived at the Tipsy Elf at a quarter past noon, they would likely be surprised, perhaps speculating that B must have had some ulterior motive for withholding information, and concluding that their exchange had been less cooperative than they'd thought. By contrast, at least in some transparently adversarial contexts, B won't be taken to have implicated anything by their utterance in (9). For example, suppose that (9) occurs in the context of A's heated deposition of B, and they both know that A's goal is to uncover evidence that B committed a crime. In this sort of context, nobody is expected to volunteer more information than strictly necessary. And so it is not tempting to understand B as implicating that they have no closely related information to offer which might be relevant to A.

3 Genre as the Question Under Discussion?

Having now laid out our four genre distinctions, we turn to the question of what unifies them. Our answer will be that they must all be understood in terms of conversation plans. But our path to that idea begins with the simpler hypothesis that a conversation's genre is determined by its Question Under Discussion, or QUD. This idea is attractive because QUD-based models of discourse have already been used to explain a wide range of other phenomena (Beaver et al., 2017), and are particularly useful in explaining why certain speech acts are more natural and expected than others. They do this by predicting which speech acts are relevant to the question that interlocutors are addressing at a given moment in conversation. So, here is a hypothesis: what it is for a speech act to be natural and expected in the ways that we've been talking about is just for it to be relevant to the QUD, and variation in conversational genre is just variation in QUD.

Ultimately, we will argue that this idea can explain the informative–practical distinction, but not the other genre distinctions that we've discussed, all of which can vary independently of QUD. But it will be useful to introduce the QUD account first, because our positive view will generalize rather than reject it.

QUD models' motivating insight is that, at any given stage, a conversation tends to be organized around both the Common Ground (CG), which is the body of information that the interlocutors are taking for granted for the purpose of the conversation, as well as the Question Under Discussion (QUD), which is the question that the interlocutors are seeking to answer. A speech act will normally be counted as relevant only if it constitutes progress toward answering the QUD, in light of the current state of the CG. In particular, an assertion is relevant only if it at least partially answers the QUD, in the sense that accepting its content into the common ground would eliminate one of the hitherto-live answers to the QUD. And it is relevant to ask a question only if, in this context, it is a subquestion of the QUD, in the sense that a complete answer to this new question would be at least a partial answer to the prior QUD.[4]

Here's our idea about how to understand the informative–practical distinction: The QUD of an informative conversation is an informative question—a question that can be resolved by acquiring new information. By contrast, the QUD of a practical conversation is a practical question, or "decision problem"—a question about what one or more of the interlocutors is to do, which has to be answered by making a decision, rather than by learning something new. Assertions are more natural and expected in informative conversations because they are speech acts for imparting information, and so are the sorts of speech acts with which we can give literal answers to informative questions. Directives, as well as practical suggestions and proposals, are more natural and expected in practical conversations, because they are speech acts for proposing actions, and so are the sorts of speech acts for which we can give literal answers to practical questions.[5]

This view can explain why declaratives tend to come off as indirect proposals, suggestions, or directives in practical conversations, whereas imperatives can sometimes to come off as mere attempts to share information in informative conversations. Recall (5), which comes off as a practical suggestion when uttered in a conversation about where to have dinner:

[4]We will stick to informal descriptions of QUD models here. For technical details, consult Roberts (2012).

[5]The idea that directives are felicitous only when they are being used to answer decision-problem QUDs has been defended by Kaufmann (2012, 159–161), Harris (2021, 1079), and Roberts (2022, §5). There are interesting questions about how to build a formal model that could help us to rigorously distinguish informative and practical questions and predict the felicity of speech acts of different kinds, but we won't attempt that here. For some previous attempts to model the interactions of decision problems and QUDs, see van Rooy (2003) and (2014).

(5) Calexico has good tacos.

Suppose that when (5) is uttered, the QUD is the question, *where should we have dinner?* We can then understand the practical reading of (5) as a relevance implicature, which arises from a mismatch between the utterance's literal content and the QUD. Since the QUD can only be answered with a proposal about what to do, the assertion that is literally indicated by (5) is not, on its own, relevant. On the assumption that the speaker is cooperative, they must have meant something additional, which *was* relevant. Since Calexico's having good tacos would be a reason to have dinner there, the best explanation of why the speaker uttered (5) is that they were trying to indirectly propose eating there—a proposal that, if accepted, would answer the QUD. This mirrors QUD theorists' explanations of other relevance implicatures (Roberts, 2012, 21).

We can give a similar explanation of the practical reading of (6):

(6) (a) Anya: I am not trying to get to Grand Central Station, but I am curious: how does one get there?
(b) Baker: Take the 5 train.

Here, Anya asks an informative question, thereby installing it as the QUD, and clarifies that she is not trying to accomplish an additional, practical goal by asking the question. To answer an informative question, one needs to provide information—paradigmatically, by means of an assertion. But Baker utters an imperative, which literally indicates that she is issuing a directive or making a practical suggestion. Taken literally, this would not be a relevant answer to the QUD. This sends Anya in search of an indirect meaning that *does* felicitously answer her question. Which relevant answer to the QUD might Baker have intended to indirectly communicate by their utterance, and why would that utterance have been a good way to give evidence that this answer is what they intended? Plausibly, Baker meant that *one gets to Grand Central Station by taking the 5 train.* Exactly how Anya works this out is an interesting question, and the details will depend on the precise theory of imperatives on offer. But suppose, for concreteness, that the meaning of an imperative specifies a plan for the addressee (Charlow, 2013; Harris, 2021). Given that the QUD is a question about how one gets to Grand Central Station, why would Baker respond with an utterance that, if taken literally, would propose a plan for Anya to take the 5 train? Plausibly, in order to make salient the proposition that this is the sort of plan that anyone could adopt in order to get to Grand Central—a proposition that *is* an answer to Anya's question.

We can give a similar explanation about why sub-clausal utterances, such as (7), are most naturally understood as assertions in informative conversations but as practical proposals in practical conversations.

(7) Calexico.

Outside of a clause whose type conventionally indicates illocutionary force, a one-word utterance is illocutionarily ambiguous. But it is well known that the QUD has broad powers to disambiguate sub-clausal answers. Compare (11) and (12):

(11) Anya: Who has the best tacos in Brooklyn?
Baker: Calexico.

(12) Anya: Where do they have mezcal around here?
Baker: Calexico.

Although Baker utters the same word in each case, she is using this word to say that Calexico has the best tacos in Brooklyn in (11), and that Calexico has mezcal and is nearby in (12). It is disambiguated in different ways by the interlocutors' assumptions that Baker is cooperatively trying to answer the QUD. Our proposal about (7) is the same: illocutionarily ambigious utterances get disambiguated as assertions when the QUD is an informative question, and as practical proposals when the QUD is a practical question.

4 From QUDs to Conversation Plans

We're happy with our QUD-based explanation of the informative–practical distinction, but it can't be extended to our other three dimensions of genre variation, because conversations vary along those dimensions even if we hold the QUD fixed. Suppose that Anya and Baker are having a conversation in which the QUD is, *Which restaurant is on the corner of 5th Avenue and 1st Street?* From this supposition alone, nothing follows about where the conversation resides on the other dimensions. Are Anya and Baker having a committal conversation in which they are seeking to offer firm answers to the QUD, or an exploratory conversation in which they are merely trying to brainstorm possible answers? Are they in a factual conversation in which they are treating the QUD as a question about reality, or a make-believe conversation in which they are treating it as a

question about some fiction or pretense? Are they addressing the question by cooperating on an answer, or having an adversarial quarrel?

This line of thought might lead one to wonder whether our four genre distinctions are too heterogenous to be explained with a single theory. But we think that conclusion would be too hasty. In order to capture the other dimensions of genre variation, we propose to think about the actual phenomenon that QUD models are models of, and to zoom out from that phenomenon to see the larger structures of which QUD models depict only parts.

These larger structures are conversation plans, which are hierarchical structures of intentions that govern and organize conversations.[6] On this view, the reason why a certain question is the QUD at a given stage of a conversation is that the participants have a shared intention to address that question at that stage.[7] But, like intentions in general, the shared intention that selects the QUD is only one element within a larger plan structure. Our other dimensions of genre variation are explained in terms of other recurring elements within conversation plans.

We'll need a theory of plans. For this, we look to the work of Michael Bratman (1987; 2014), who takes plans to be complexes of intentions, arranged into hierarchical structures that reflect the practical-reasoning process by which they were created. The plan-building process starts with an intention to do something—say, host a dinner party. This intention is *partial*, in that it represents a goal without specifying the details of how to accomplish it. In order to work out these details, the agent must choose *subplans* of their intention—further intentions that fill in the details of the plan. For example, in order to work out the details of how to host a dinner party, the agent has to choose when it will happen, whom to invite, what to serve, etc. If they decide to serve pasta, this new subplan still leaves many questions unanswered: Where will they get the ingredients, how will they cook them, and so on? They will thus have to iterate this reasoning process a number of times, resulting in a complex, hierarchical plan that connects the abstract intention with which they began, via a network of subplans, to intentions specific enough to be turned into bodily movements.

A rational planning agent forms intentions that are consistent with their other intentions, with their beliefs about what they can do, and with the need to

[6]We are inspired by other theorists who have understood conversations as jointly planned activities, although none of these theorists have spelled out the details of our account of conversation plans: Bratman (2014); Clark (1996); Grice (1975); Grosz and Sidner (1986); Roberts (2012).

[7]See Friedman (2025) for a recent argument that individual inquiry should be understood in terms of intentions to address questions, in a way that is congenial to our argument here.

constructively and coherently flesh out their existing plans. On Bratman's view, it is part of the nature of intentions not only that they function as elements in larger plans, but also that our mechanisms for forming them are responsive to these sorts of rational pressure.[8]

This sort of planning is useful for ambitious but cognitively limited creatures, like us. It allows us to break complex, multivariate decisions into tractable chunks. Because the outcomes of earlier decisions constrain the options from which we choose when making later choices, these later choices are easier to make. For both of these reasons, hierarchical planning allows us to accomplish more complex actions than would otherwise be allowed by our limited cognitive resources, and to pursue more abstract goals, in the sense that they would have to be pursued by very different specific actions in different circumstances.

Bratman argues that shared plans work in largely the same way as individual plans. The main new ingredient is for a group of agents to have a shared intention, which, on Bratman's view, requires that each intends to do something together with the others, intends to pursue this end by means of meshing subplans, and is aware that the others have these intentions.[9] A shared intention has meshing subplans when there are intersubjectively coherent intentions aimed at its satisfaction which, taken together, add upp to a complete representation of a way of pursuing the intention. This is the main idea we take from Bratman. For example, if Tim and Ruth have a shared intention to host a dinner party, then each of their plans about their own contributions must fit together with the other's. If Ruth intends to handle the logistics and expects Tim to cook, then Tim should intend to cook and expect Ruth to handle the logistics.

A more detailed example is illustrated in Figure 1. Several features of this plan bear attention. First, notice that the shared intentions have unshared intentions both as subplans and as superplans. Shared intentions tend to have unshared subplans because it is ultimately individuals who must translate the shared plans into particular bodily movements, such as picking up the telephone to call Pete. This is part of why shared intentions need to have meshing subplans. On the other hand, agents typically engage in joint activity for their own reasons, and so shared intentions are themselves typically subplans of agents' unshared intentions. Whereas Tim wants to have a dinner party as a way of pursuing his intention to learn about Steve's vacation, for example, Ruth's main

[8]For proposals about exactly how these rational requirements should be formulated, see, e.g., Bratman (1987); Broome (2013), and Worsnip (2021).

[9]Bratman's theory of shared intentions is more complex and demanding this, in ways that we don't think are relevant. For details, see Bratman (1992; 2014).

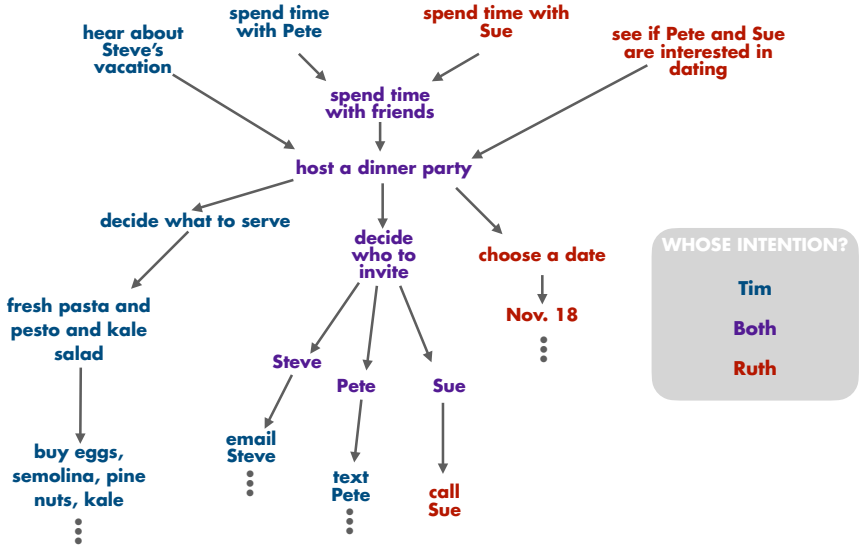

Figure 1: An example joint plan. Arrows point from intentions to their subplans.

aim is to spend time with Sue. Ultimately, we will argue that conversation plans have structures like this as well.

A crucial step toward our theory of conversational genre is to understand QUD models as models of shared intentions that serve as elements in larger shared plans. One thing that this interpretation has going for it is that it captures the intuitive motivation of QUD models, which Roberts introduces by saying that "[d]iscourse is organized around a series of conversational goals and the plans or strategies which conversational participants develop to achieve them" (Roberts, 2012, 3). She describes QUDs as "discourse goals" (Roberts, 2012, 26), distinguishing these from "domain goals," which are "things we want to achieve quite apart from inquiry" (Roberts, 2012, 7). But, in later refinements of her model, Roberts conceives of discourse goals as a "distinguished type of domain goal, those the interlocutors are jointly committed to achieving in the discourse itself" (Roberts, 2018, 323). And, even more recently, she describes domain goals as "an organized body of the interlocutors' goals, plans, and priorities" and explains that her framework "borrows directly from the insights of planning theory (Bratman 1987)" (Roberts, 2022).

More importantly, the idea that QUDs model shared intentions can help to explain the relationship between QUDs and relevance. This explanation relies on Grice's (1957; 1969) idea that what a speaker means—and the content and force of

the speech act that they perform—is a matter of their communicative intention.[10] Given this assumption, we can say that for a communicative act to be relevant to the QUD is just for the speaker's communicative intention to be a coherent and constructive subplan of the interlocutors' shared plan to address the QUD. Suppose that Baker and Anya have a shared intention to answer the question, *Do they have horchata at Calexico?* In this context, if Anya knows the answer, it makes sense for her to pursue their common goal by forming a communicative intention to convey this information. The fact that her communicative intention is a constructive and coherent subplan of their shared intention to answer the QUD is what makes Anya's assertion relevant. By contrast, if she were to form a communicative intention to assert that Wes Anderson's first film was Bottle Rocket, her intention would *not* be a constructive subplan of their shared plan, which is what would make her assertion *ir*relevant. We can explain the (ir)relevance of questions in a similar way. To ask a question, according to QUD models, is to propose a new QUD, which ought to be a subquestion of the prior QUD. On our interpretation, proposing a new QUD is proposing the adoption of a shared plan to answer a question. In general, since answering a subquestion of q is a useful step toward answering q, a shared intention to answer a subquestion of q is a coherent and constructive subplan of a shared intention to answer q.

On this interpretation of the QUD model, Roberts' notion of relevance turns out to be just one instance of the broader phenomenon of plan coherence.[11] To make a relevant move in a conversation is just to do something with a communicative intention that is a coherent and constructive subplan of the interlocutors' shared plans. There is considerable independent motivation for idea that we strive to form coherent plans, and that we are sensitive to others' failures to do so, particularly when we are engaged in joint planning, and so this interpreta-

[10]Specifically, a communicative intention is an intention to produce a certain mental state in one's addressee(s), in part by revealing the intention to do so. We won't explicitly argue for a Gricean account of speech acts here, though we take the dependence of our theory of genre on a Gricean account to be a consideration in the latter's favor. For a comparison to other theories of speech acts, see Harris et al. (2018). For defenses of Gricean theories of speech acts, see Strawson (1964), Schiffer (1972), Bach and Harnish (1979), Harris (2014; 2019; 2020), and Unnsteinsson (2017; 2022; 2023).

[11]Here we are inspired by Vesga and Starr (2025), who have proposed a similar generalization, as well as by Clark (1979) and Hawkins et al (2015), who present experimental evidence that we are sensitive to interlocutors' background intentions in the ways that we formulate and answering questions. See also Friedman (2025), who argues that inquiry in general is guided by question-directed intentions, and that the norms governing it are just those of practical rationality.

tion of QUD-relevance grounds it in a more general and independently motivated notion.

Our strongest argument for the idea that QUD models are representations of shared intentions will emerge over the next several sections. In brief, the argument is that this interpretation allows us to generalize QUD-based explanations to explain how other dimensions of genre variation influence which kinds of speech acts are natural and expected. Although the other dimensions can't be explained in terms of the QUD itself, because they can vary independently of QUD, they can be understood in terms of other elements in our shared conversation plans, which interact with the QUD in the same way that different intentions within complex plans typically interact. This gives us a powerful reason to think of the QUD as a model of just one element within larger planning structures.

5 Exploratory Conversation and Attentional Moves

With a sketch of conversation plans on the table, we'll now flesh it out with our accounts of the other genre distinctions. We begin with the distinction between committal and exploratory conversations.

Consider again the example of doctors gathered around a white board, brainstorming possible causes of a patient's diagnosis:

(8) Thirteen: A brain tumor is restricting blood flow to her hypothalamus.
Kutner: She fell and hit her head.
Taub: She has lupus.
Thirteen: A toxin.

There are good reasons to think that the QUD in this conversation is just the question:

(13) *What is causing the patient's symptoms?*

The doctors are having their conversation in order to answer this question. If they had the answer, they would move on to other business. At a later stage in the conversation, for which this stage is the preparation, they will try to answer the question by ruling out some of the possibilities they have generated and explored. Even at this stage, their utterances are felicitous only if the suggested answers really are live possible answers to the question. If one of the doctors were to suggest that the patient is very sleepy, or that their symptoms are caused by a

microchip implanted in them by John F. Kennedy, for example, the others would likely treat these utterances as infelicitous, because they are all ruling out these hypotheses, and would prefer to keep it that way. In this context, whether an utterance is relevant isn't a matter of whether its content is incompatible with one of the live answers to the question, (13); rather, an utterance's relevance comes down to whether its content *is* (or should be) one of the answers to the QUD that the participants are treating as live. This is why it's okay for the doctors to make utterances with inconsistent contents, as long as they are happy to treat those contents as competing, but live possibilities.

The need for this kind of conversation tends to be idealized away in Robert's discussion of the QUD model, where questions are thought of as establishing partitions of complete answers at the point of utterance (e.g., Roberts 2012, 5). It sometimes sounds as though a newly installed QUD automatically makes salient the set of live possible answers, which the interlocutors are then free to survey and choose between. But, in real life, there is considerable cognitive distance between entertaining *wh*-question like (13) and being able to recognize, entertain, and call to mind the various answers that are worth considering.

One problem is that you can be curious about a question without knowing some or all of its possible answers. Friedman (2013) calls this "hypothesis ignorance" or "abductive ignorance." Ignorance of this kind is particularly common when we face explanatory questions, such as *Why is the sky blue?*—or, we would add, *What is the cause of this patient's symptoms?* Friedman argues that, for someone who is curious about a question, but abductively ignorant, the next stage of rational inquiry is to generate hypotheses. A second problem is that even agents who know a possible answer might have trouble calling it to mind in response to a particular question.[12] Someone might know that Regina is the capital of Saskatchewan but tend to forget when asked to name Canada's provincial capitals, for example. In contexts of collective brainstorming, there may be possible answers that everyone knows, without everyone knowing that everyone knows. Finally, in some contexts, one or more of the participants in a conversation may have mistakenly ruled out a possible answer. Exploratory conversation is useful for rectifying all of these barriers to the understanding and resolution of questions. It is a collective activity in which we generate, draw attention to, and share possible answers to a question, ultimately in the service of answering it.

How should we characterize the individual speech acts by means of which we accomplish these aims—the kind that the doctors are performing in (13)? Here,

[12]Cf. Hoek and Bradley (2022) on the role of questions in collective deliberation.

we take inspiration from a number of theorists who have argued that sentences containing epistemic possibility modals, such as (14), are at least sometimes used to draw attention to live epistemic possibilities without eliminating other live possibilities—speech acts that we will call "attentional moves."[13]

(14) She might have lupus.

What might someone be trying to accomplish by performing an attentional move? Just the sorts of things that we try to accomplish in exploratory conversations: a doctor might utter (14) in order to draw attention to a possible cause of the patient's symptoms that their interlocutors didn't know about, one that they did know about but aren't actively considering, one that everyone knew about but wasn't sure if the others recognized, or one that the others might have mistakenly ruled out.

Our hypothesis is that exploratory conversations are those in which the participants have a shared plan to address the QUD only with attentional moves, at least temporarily. In many cases, this intention about how to address the QUD will be a subplan of the QUD intention. This will be true in cases like the doctors' conversation, where brainstorming about possible answers to a question is a preliminary step toward answering it, designed to generate as many candidate answers as possible before eliminating some.[14] This is a bit like Fisher and Ury's (1981) justification for their recommendation to begin negotiations with a brainstorming phase before moving on to bargaining.

We can also imagine purely committal conversations, in which the participants have both a shared intention to get down to business and answer the QUD rather than merely generating possible answers. Again, the latter intention will often be a subplan of the former. This kind of situation could arise if answering the QUD is particularly urgent, or if the participants feel that they have spent enough time exploring answers. Suppose that the patient's worsening condition

[13]See, for example, Veltman 1996, Yalcin 2007, Ciardelli, Groenendijk, and Roelofsen 2009, Dever and Schiller 2020. For some experimental evidence that epistemic possibility modals are used to perform attentional moves, see Bade et al. (2022). We won't weigh in on the debates between these theorists about how to formally model attentional moves, or about the division of labor between semantics and pragmatics that makes epistemic modals particularly useful for performing them.

[14]We can also imagine a group of people who first decide to brainstorm—say, in order to practice their brainstorming skills—and then choose a question to brainstorm about as a means to that end. In this case, their QUD-intention would be a subplan of the shared intention to engage only in attentional moves. We think that cases like this are pretty rare in real life.

demands immediate treatment, and our doctors have to decide between the possibilities they've considered right away, for example. In this context, we need an answer to the question, not another merely possible answer.

Many conversations are neither strictly committal or strictly exploratory, in that participants have no shared intention to restrict themselves to either assertions or attentional moves. In these conversations, speakers typically seem to mark attentional moves, either with epistemic possibility modals or in some other way. And when they utter a bare declarative without such marking, the best explanation is that they are performing an assertion. In this sense, we might think of assertions as the default use of bare declaratives. Our explanation of the fact that exploratory conversations sometimes license attentional readings of bare declaratives is that this assertoric default can be overridden by the participants' shared intention to restrict themselves to attentional moves, for it leads them to expect attentional moves rather than assertions.[15]

We can think of these expectations as resulting from a kind of relevance constraint that generalizes Roberts' (2012) theory of QUD-relevance. In addition to being relevant to the QUD, a speech act should be relevant to all of the shared intentions that make up the interlocutors' conversation plan. Given our interpretation of relevance as subplan coherence, this is what we should expect: just as a speaker is under rational pressure to make their communicative intention a coherent and constructive subplan of the QUD-intention, they should also make it a coherent and constructive subplan of their other shared intentions. QUD relevance is an instance of conversation-plan relevance, which, in turn, is an instance of plan coherence.

6 Make-Believe and Presupposition Plans

Next, consider the difference between factual and make-believe conversations. Take the following exchange:

[15]Why do declaratives seem to default to assertoric uses, and what does overriding this default entail? Possibly, the meaning of a declarative makes it especially well suited for assertions, and attentional uses of bare declaratives are indirect speech acts, in which the speakers make as if to assert a proposition in order to indirectly draw attention to it. On this view, the implicature could be triggered in conversations like (13) by the mismatch between the sentence's meaning and the interlocutors' shared intentions. Alternatively, the meaning of a declarative could be a proposition, and that whether a speaker who utters it is thereby asserting or drawing attention to it is a pragmatic matter. On this view, the default would be a matter of participants' expectations about one anothers' intentions. We won't try to decide between these two options here.

(15) Anya: How many eyes does Gilitrutt have?
Baker: Three.

Uttered within a make-believe conversation, Anya won't come away believing that Gilitrutt exists, or has three eyes. Rather, Baker's aim is for Anya to imagine, or make-believe, that Gilitrutt has three eyes. In this scenario, it would miss the broader purpose of their conversation for Anya to respond by pointing out that what he says isn't true.

On the other hand, if Anya thinks that Gilitrutt is real, then the same exchange could happen within a factual conversation. In this case, it would be perfectly natural for Anya to ask Baker how he knows, or to express doubt about the accuracy of Bakers's claim. And in this case, Baker's claim will be fully successful only if Anya ends up believing what Baker has said.

Notice, however, that the QUD is the same in both cases, imposing the same relevance constraints on subsequent conversational moves. What differs between these two cases is not the QUD, but the common ground. In the factual version of the conversation, Anya and Baker are treating as common ground only propositions that they believe and for which they have good evidence. (In other words, they are trying to treat the common ground as the body of propositions that they know to be true.) By contrast, in the make-believe version, they are working with a different sort of body of information as common ground—one that includes some propositions that they are treating as true only for the purpose of their game of make believe.[16]

We can say something similar about Alvy and Allison, from Section 2.3, who suppose that a second assassin helped to kill JFK:

[16] As noted in Footnote 3, we have been assuming a theory of fictional discourse on which the contents of fictional claims can be the same as factual claims, but interlocutors bear different attitudes toward those contents (e.g., make-believe vs. belief). Some theorists have instead held that fictional discourse involves ordinary beliefs about fictional contents, which differ from factual contents in some way—for example, because they include (possibly covert) fiction operators. If a content-based theory is true, then our explanation of the fictional–factual distinction isn't right. For example, it will turn out that Anya is asking a (covertly) different question than she would have asked if (15) had been part of a factual conversation. The question would be about how many eyes Gilitrutt has *in the fiction*, for example, and Baker could be asserting that Gilitrutt has three eyes *in the fiction.* It would be easy enough for our broader theory to make sense of the factual–fictional distinction, given a theory of this kind. The factual–fictional distinction would turn out to be an example of genre variation that results from differences in the QUD. Roughly: make-believe conversations are those in which the QUD is a question about how things are according to the fiction.

(16) Alvy: Who was in on the conspiracy?
Allison: Everyone was in on it: Earl Warren, Lyndon Johnson, the CIA.

In this context (unlike one involving real conspiracy theorists), Allison isn't expressing a belief, or trying to get Alvy to believe what she says, and Alvy wasn't (necessarily) seeking factual information when he asked his question. Rather, Alvy's question, installed as the QUD, guides the conversation by setting up some modification of a common ground that includes suppositions that neither Alvy nor Allison takes to be true, but which they are merely treating as true for the purpose of the conversation, which is (at least in part) to figure out what would follow from these suppositions.

Common-ground and QUD models of conversation have been designed, from the start, to make sense of the fact that we sometimes take different kinds of attitudes toward the propositions in the common ground. For example, Stalnaker defines common ground as the set of propositions that the participants commonly believe that they "accept (for the purposes of the conversation)" (2002, 716; see also 1984; 2014). He explains the role of acceptance as follows (1984, 79):

> acceptance…is a generic propositional attitude concept with such notions as presupposing, presuming, postulating, positing, assuming and supposing as well as believing falling under it. To accept a proposition 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.

Yalcin (2018, 405) has pointed out that common-ground models have one fairly surprising feature. Strikingly, they can make compelling predictions about the role of shared information in conversation, while staying silent on which specific attitudes the interlocutors take to that information. Stokke (2023) makes a similar point about fictional discourse, arguing that since they are pragmatically similar to factual discourse, the contexts probably have similar structure. As our two interpretations of (15) illustrate, the same goes for QUD models: we can ask the same question, thereby installing the same QUD, whether we do so in the pursuit of factual knowledge, the next element in a fictional narrative, or the next step in a *reductio* proof.

Yalcin refers to the attitude that interlocutors collectively adopt toward the common ground—be it knowledge, imagination, or provisional supposition—as

the "conversational tone" (2007, 1008). Our proposal in this section, in effect, is that the distinction between factual and make-believe conversations is a matter of a conversation's tone.

This isn't the full story, though, because it leaves out an account of what makes it the case that a conversation has one or another tone. In virtue of what does a group of interlocutors treat common ground as a body of shared knowledge on one occasion, but a collection of bizarre fictions on another occasion? Predictably, our answer is to posit yet another kind of recurring element within conversation plans—a shared intention to treat one or another body of information as common ground, which we will call the CG-intention. In a way, this idea is already there in Stalnaker's reference to "the purpose of the conversation" in his analysis of common ground. In a brief aside in a more recent work, Stalnaker spells out the point further: "In simple straightforward serious conversations, what is accepted will coincide with what is believed, but some conversations will involve some mutually recognized pretense about what is believed, or some tacit suppositions that are made in order to further the mutual aims of the conversation" (Stalnaker, 2014, 4). We propose to take this claim about "mutual aims of the conversation" seriously, and incorporate it into our theory of conversation plans.

Often, we choose what sort of common ground to work with first, and then adopt QUDs as subplans: we ask questions as "strategies of inquiry" to pursue the facts (Roberts, 2012, 6–7), or as a means to the end of finding out or deciding what some shared fiction or pretense is like. But things can also be the other way around: we sometimes temporarily treat non-factual information as common ground as a means to the end of answering a factual question, as in reductio arguments, conditional proofs, or when counterfactual possibilities have a bearing on what we should think about the actual world.[17] One benefit of the idea of conversation plans is that it allows us to make sense of these different kinds of instrumental relationships that QUDs can bear to our choice of what to treat as common ground.

[17]What we say here is related to the idea that we sometimes operate temporarily with "local" or "derived" contexts, sometimes in ways that are signaled linguistically, but sometimes not. See Roberts (1989) and Stalnaker (2014, §4.2).

7 Coherence and Conflict in Unshared Plans

Finally, consider the distinction between cooperative and adversarial conversations. Cooperation is presumably a stew with many ingredients, and will tend to drop off when, for example, agents are in certain kinds of emotional states, don't trust each other, and lack shared pro-social dispositions and background information. We don't claim to make sense of all of these aspects of cooperativity. Our claim here is only that one important component of cooperativity in conversation involves the degrees of both coherence and transparency in the intentions that interlocutors are seeking to satisfy by by means of conversation, and of which the shared intentions that make up their conversation plans are therefore subplans. Insofar as interlocutors are aware of and inclined to try to support each other's intentions, the conversation will tend toward cooperativity. If the intentions relevantly conflict, the conversation will tend to be more adversarial. And if interlocutors simply aren't aware of each other's intentions, their ability to cooperate will be impaired (see figure 2).

This goes for joint activities in general, and not just conversations. Recall Ruth and Tim from §4, who are planning a dinner party. Suppose that Ruth is planning the party because she intends to introduce two friends whom she thinks would make a good couple, but Tim thinks that this would be a terrible idea, and intends to prevent these friends from meeting. This conflict might not impede some of the decisions they face, such as those concerning the party's time and location, but when they have to decide whom to invite, it could lead to an impasse, as Ruth and Tim will struggle to find shared subplans of their conflicting private intentions. When this source of adversariality is taken to extremes, it may be impossible for agents to act together at all. But, in less extreme cases, conflicting intentions will merely constrain and add friction to the process of forming shared plans.

Each agent may intend to achieve a whole range of goals in a conversation, any one of which may be compatible or incompatible with any one of their interlocutors' intentions. Moreover, the degree of compatibility of some pairs of intentions may be more relevant than that of some other pairs. The cooperative–adversarial distinction is therefore graded and multidimensional: adversariality will tend to increase with both the number of conflicts and with their degree of relevance to the interlocutors' shared plans.

Consider again these examples:

(9) A: Were you at the Tipsy Elf at noon yesterday?

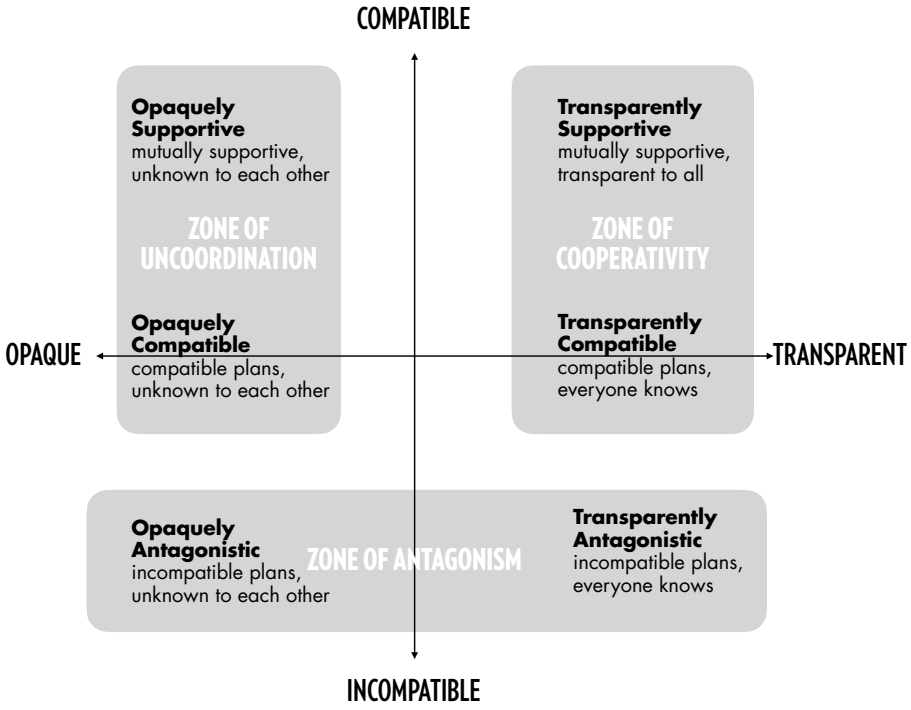

Figure 2: Two dimensions along which two interlocutors' intentions can vary. They can range from being mutually supportive to incompatible, on one hand, and from transparent (in the sense of being mutually known) to opaque, on the other.

B: No.

(10) A: Were you at the Tipsy Elf at noon yesterday?
B: No. But I was there at half past noon.

We pointed out that in many cooperative conversations, but not in at least some adversarial conversations, B would be understood, in (9), to be implicating that they hadn't been at the Tipsy Elf at any time close to noon.

We can now give a more detailed explanation, which depends on the idea that interlocutors reason not only about what is relevant to their shared plans, such as their shared intention to address the QUD, but also about the unshared plans of which those shared intentions are subplans. A is asking about B's whereabouts at noon yesterday, and thereby proposing a new shared QUD-intention, *for a reason*—i.e., in order to further some antecedent plan. Suppose that B knows, or

can infer, what this antecedent plan is, and that it would be served by information other than an answer to the QUD. For example, perhaps B can infer that A asked their question as a way of finding out about B's day more generally, or as a way of finding out about whether B bumped into a mutual friend who had also been at the Tipsy Elf. Even if B can't infer A's specific motive, B might guess that it is likely to be some intention that could be furthered by the information that they were at the Tipsy Elf at some time close to noon. In these scenarios, B recognizes that they could cooperatively further A's plans by sharing extra information, beyond an answer to the QUD, as they do in (10). And, if all of this background is sufficiently obvious to A and B, but B does not share any additional information that would be relevant to B's plans, as in (9), B might be taken to implicate that they have no information that would be relevant. The inference by which this implicature is calculated is straightforward: if B is cooperative and had relevant information, they would have offered it, and so they must not have any.

This reasoning pathway is similar to other relevance-based reasoning; it's natural to reconstruct the case by saying that B implicated that they had no further relevant information to add, because if they did, they would have offered it. But notice that in this case, we're not talking about relevance to the QUD, as in the cases that we discussed in §3, or even relevance to shared plans more generally, as in §5. In this case, it's what's relevant to A's unshared plans that guides the pair's reasoning. In §5, we suggested that the concept of relevance that plays a role in pragmatic reasoning can be reduced to plan coherence: we judge speech acts relevant insofar as the communicative intention behind them is a constructive and coherent subplan of the interlocutors' shared plans. But as we've just seen, cooperative speakers seek to further not only their shared plans, but also their interlocutors' unshared plans, insofar as they can infer what those plans are. If we are right, then what Grice called "the accepted purpose of the talk exchange" can be quite complex.

Things are very different if A and B have intentions that conflict in relevant ways. Suppose, for example, that A is interrogating B with the aim of establishing that B committed a crime at the Tipsy Elf yesterday, B intends to maintain their innocence, and each knows of the other's intentions. In that case, it would be irrational for B to offer anything beyond a simple "No," as in (9), and A would be unlikely to understand them as implicating that they weren't there at any nearby time.[18] In the cooperative scenario, B wanted to help further the intentions that

[18]One source of indirect support for what we say here comes from the literature on the effects of trial consultation and training on court witnesses. For example, Boccaccini et al. (2005) found

lay behind A's question, and so the best explanation of why they didn't volunteer extra information was that they didn't have any. But in the uncooperative case, both A and B know that B is motivated *not* to help further A's plans, and so a good explanation of their silence is that they're simply withholding information that would further A's aims, but not B's. Therefore, we get no implicature.

We can generalize from this case to the following principle:

(17) **Plan Relevance**
A speech act is conversationally relevant only to the extent that the speaker's communicative intention is a coherent and constructive subplan of either (i) the interlocutors' shared plans, or (ii) unshared plans of the addressee(s) that are compatible with the speaker's own plans and that the speaker is aware of and disposed to further.

We see this as a generalization of Roberts' concept of QUD relevance, and as a precisification of Grice's broader idea of what it is for a contribution to be relevant to "the accepted purpose or direction of the talk exchange" (Grice, 1989, 26). Specifically, we think that it goes some way toward clarifying what the "purposes" of a conversation are, and which purposes matter.

This theory helps to clarify the special role that cooperativity has often been held to play in human communication. Transparently cooperative conversations, built around rich networks of mutually understood and mutually supportive plans, give us ample resources for inferring our interlocutors' communicative intentions and making our own communicative intentions known. This allows us to communicate more things more efficiently than we otherwise could. By contrast, when interlocutors enter conversations with conflicting intentions, this limits the kinds of conversation plans that they can construct, with the result that they have to work harder to be understood, and sometimes with the result that indirect communication requires more effort, making communication less efficient in a concrete way.

With this discussion of unshared superplans of shared intentions out of the way, we can now offer a more complete account of the structure of conversation plans. See Figure 3 for a schematic example.

The genre-defining elements in a conversation plan, from top to bottom in Figure 3, are (i) agents' unshared intentions about what to do in the conversation, which determine whether the conversation will be cooperative or adversarial, (ii)

a resulting decrease in witnesses' expressiveness and a corresponding increase in likelihood of acquittal by jury.

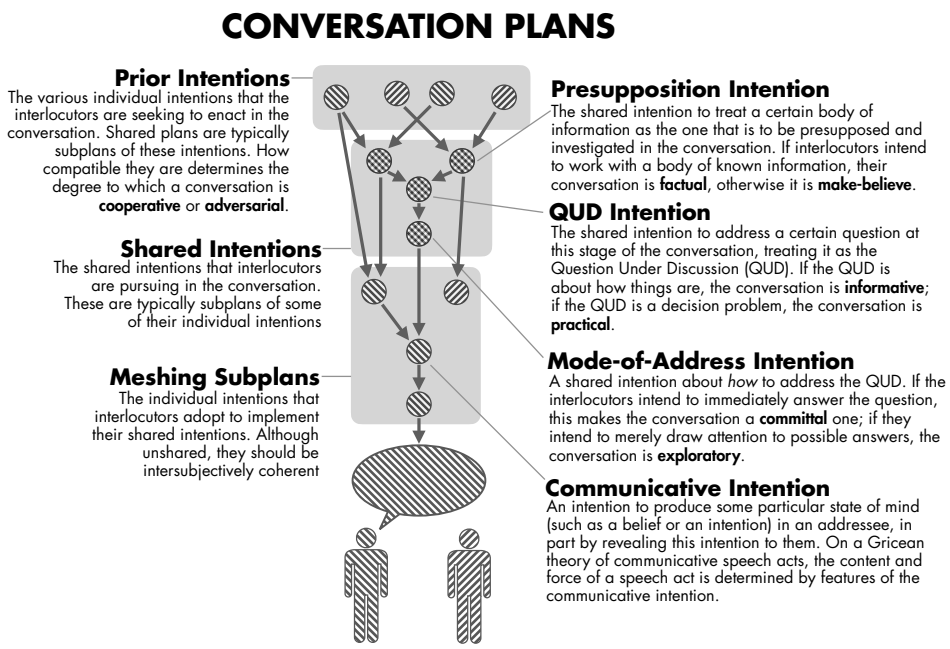

Figure 3: An schematic example of a conversation plan, including (inter alia) each of the kinds of elements that we've discussed here.

the presupposition intention to take one or another attitude toward the presupposed information, which determines whether the conversation will be factual or make-believe, (iii) the QUD-intention to address a certain question, which can make the conversation either informative or practical, (iv) the mode-of-address intention about how to address the QUD, which can make the conversation either committal or exploratory, and (v) the speaker's communicative intention, which fixes the content and force of their communicative act, and which must (in order to be rationally cooperative) be a constructive and coherent extension of the rest of the plan.[19]

[19]In Figure 3, we have represented the mode-of-address intention as a subplan of the QUD intention, which is in turn a subplan of the presupposition intention, reflecting the idea that we normally choose whether to answer or brainstorm about the QUD after settling on which QUD to address, and we normally address a QUD in the service of a plan of either discovering the facts or engaging in make-believe. However, we can imagine each of these subplan relations reversed in unusual cases.

8 Genre and Grice's Maxims

We claimed in §2 that some standard models and tools in pragmatics have been designed in ways that apply only to certain conversational genres. Grice's maxims of conversation are a good example, as Grice himself recognized: "I have stated my maxims as if this purpose were a maximally effective exchange of information; this specification is, of course, too narrow, and the scheme needs to be generalized to allow for such general purposes as influencing or directing the actions of others" (Grice, 1989, 28). In light of what we've done in this paper, we can interpret Grice as follows. On one hand, the cooperative principle itself can be understood as an almost purely structural constraint on conversation plans:

> Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged. (Grice, 1989, 26)

This is just a Bratman-esque requirement to adopt and act on coherent and constructive subplans of our shared intentions, together with only the slightest substantive assumption—namely, that we have a shared plan to have a conversation that is organized around some goal or other. The cooperative principle itself prescinds from any assumptions about the nature of these goals. By contrast, Grice's formulation of the maxim of quality presupposes a certain kind of conversation with certain kinds of goals:

> Try to make your contribution one that is true…:
> 1. Do not say what you believe to be false.
> 2. Do not say that for which you lack adequate evidence. (Grice, 1989, 27)

It doesn't make sense to describe directives and practical suggestions as true (or false), or to assess the evidence for them, and so Grice's maxim of quality doesn't apply within practical conversations. Likewise, it is not a requirement within brainstorming conversations that the possibilities we raise be true, as long as they are worthy of consideration. And it is definitional of make-believe conversations, as we have understood them, that speakers presuppose and assert what they do not believe to be true. Grice's maxim of quality is therefore operative only in informative, committal, and factual conversations, in part because

the conversation plans that organize these conversations don't prioritize the exchange of true and justified information.[20]

Unlike quality, Grice relativizes the maxim of quantity to "the current purposes of the exchange" (1989, 26–27), making it flexible enough to apply across genres. (We needn't violate the maxim of quality by sharing very little information in conversations where we have a shared plan to demonstrate riddles, for example.) The maxim of manner enjoins us to "avoid ambiguity," "obscurity of expression," and "unnecessary prolixity," and to "be orderly"—all good things to do if our aim is to efficiently share information, but less so if we are trying to satirize a long-winded and ineloquent colleague. Grice's formulation of the maxim of relation ("Be relevant.") is unhelpfully vague, and we can perhaps think of it as implicitly relativized to conversational goals. However, as we have shown, the most influential attempt to sharpen it up (Roberts' theory of QUD-relevance) made it more genre specific, so that it had to be modified to apply to exploratory conversations.

On our interpretation, the Cooperative Principle is just the general, structural rational requirement to build and act on coherent conversation plans. The maxims attempt to spell out particular consequences of this requirement, sometimes by making substantive assumptions about the contents of conversation plans (and therefore about conversational genre).

9 Genre and Pragmatic Competence

Why do conversational genres exist? What do they do for us?

A genred conversation is governed by a conversation plan that provides participants with rich information about one another's communicative intentions. Conversation plans can do this because communicators are under cooperative-rational pressure to form communicative intentions that are coherent and constructive subplans of the intentions that make up their conversation plans. If I can safely assume that your next speech act will be a committal answer to a certain informative QUD, that we are working with a factual common ground, and that you are doing your best to cooperatively pursue various private ends that we each have, then I can infer a great deal about what the content and force of

[20] In a similar vein, we are tempted to interpret Williamson's (2000, ch.11) knowledge norm of assertion, and the various tweaks proposed by others, as an instrumental requirement arising only in conversations where interlocutors have a shared intention to exchange knowledge (cf. Carter (2025)). However, we don't have space to offer a detailed defense of this view here.

your speech act is likely to be before you even open your mouth. By designing your utterance in light of the conversation's genre, you can reasonably intend to communicate more by uttering less, relying on the conversation plan to help disambiguate. Conversational genres make us more efficient communicators.

Of course, we don't *need* conversation plans in order to communicate. As Buchanan and Schiller have illustrated, it's possible to assert that it's hot even when everyone already knew what the weather was like, and nobody had the goal of finding out more (2022, 74). Even when there is a conversation plan, successful communication can run counter to it. Buchanan and Schiller (ibid.) imagine someone replying to a question about their uncle by saying that their food is cold, and rightly argue that although this assertion would be uncooperative, it could still be understood, which would amount to successful communication. They conclude that the fact that a speech act "was *relevant*, or *informative*, or *interesting*, or *pertaining to a mutually understood topic*" can at best be features that defeasibly guide the audience's interpretation; only the speaker's intentions can make it the case that they mean one thing as opposed to another. We agree, and we would say the same about the defeasibility of conversation plans in general. Indeed, this follows from our view, together with the fact that people sometimes form intentions, including communicative intentions, that don't coherently and constructively flesh out the plans they already have, either because they don't yet have any other relevant plans, or because the new intention conflicts with their prior plans. (For an extreme example, think of someone who makes a run for it on their wedding day.)

Where we disagree with Buchanan and Schiller is in the "particularist" conclusions they draw about the pragmatic theory:

> Reflecting on the defeasible, and variable, role that features such as relevance play in utterance interpretation should make us suspicious of the search for robust, explanatory, predictive principles that offer much by way of improvement on the Gricean platitudes concerning what, as Grice himself once put it, 'a decent chap should be expected to do.' " (Buchanan and Schiller, 2022, 77)

We think that Buchanan and Schiller reach this overly pessimistic conclusion by expecting the wrong kind of answer to the central question of pragmatics, which is, "How do humans communicate?" To see what we mean, consider two different ways that someone could answer a different question:

(18) How do you get from New York City to Austin, Texas?

One way to answer would be to give necessary-and-sufficient conditions that would pick out all and only the possible ways of traveling between the two places. Since there are many weird, but possible ways to make the trip (e.g., walking backwards, being pulled behind a series of emus, etc.), these necessary-and-sufficient conditions would have to omit details about specific ways to make the trip. But normally, someone who asked (18) would be asking for a description of a good way that people usually get from city to city. We think that Buchanan and Schiller's pragmatic particularism, and its accompanying pessimism about the possibility of robust pragmatic theory, make sense only if pragmatic theories have to give the first kind of answer to the question of how humans communicate. If the goal of pragmatic theory were to describe how human communication works in every possible instance, then it's true that we can't give many details, beyond the fact that people recognize each other's intentions by all sorts of methods. Indeed, if pragmatic theories need to generalize across *all* instances of human communication, it's not even clear that we should be talking about communicative intentions or inference to the best explanation at all, since they may be part of just one particularly efficient strategy for communicating, rather than a necessary ingredient in every form of communication that humans engage in (Armstrong, 2021; Harris, 2025; Moore, 2018).

We think that most pragmatic theories, including the one that we have articulated, are best construed as attempts to give the second kind of answer to the question of how humans communicate—one that is akin to saying that, in order to get from New York to Austin, you fly from either JFK or EWR to AUS. Pragmatics is in the business of trying to explain how people *normally* manage to communicate so efficiently, in part because that's often a *good* way to do it. Conversational genre and the conversation plans that constitute it are part of an answer to this question not because they are the only way to communicate, but because they are part of a recurring strategy to make communicating easier, by creating a rich superstructure within which more ambitious communicative intentions can be made scrutable to others.

We think of this theory of conversational genre is part of a larger theory of pragmatic competence, which aims to account for the extra-grammatical psychological capacities and mechanisms that allow humans to communicate with each other. Assuming that humans often communicate by revealing their addressee-directed intentions, as Grice theorized, a theory of pragmatic competence will tell us about the psychological capacities and mechanisms by means of which

we form, reveal, and recognize communicative intentions.[21] Some aspects of our pragmatic competence may also be required in order to fully take advantage of the grammatical affordances of natural language. A plausible example is the way in which a capacity for mindreading is required for the fully competent uses of attitude verbs, epistemic modals, and evidentials. Our capacity to navigate conversation plans provides us with other cases, since natural language includes mechanisms that interface with conversation plans. One example, discussed by Roberts (2012), is prosodic focus as a device for probing and reinforcing QUDs. Another plausible example might be the use of the verb, 'suppose,' as a device for proposing a plan to operate with a non-factual common ground.

10 Conclusions and Open Questions

We have argued that many conversations belong to recognizable genres, and that genre distinctions reduce to different dimensions of variation within conversation plans. These plans put pressure on rational, cooperative communicators to form certain kinds of communicative intentions, thereby making these intentions more predictable and scrutable, and communication more efficient. Although people sometimes perform contrary-to-genre speech acts, and can thereby communicate in genre-atypical ways, our theory predicts that this will typically require communicators to override genre-based expectations, and that doing this risks both uncooperativity and inefficiency.

Conversational genre is a relatively unexplored topic, and one that raises plenty of questions that we lack the space to address. There are other genres worth considering, such as negotiation, interrogation, joking around, "bull sessions" (Frankfurt 1986). We would like to think more about the relationship between the sort of non-conventional genres that we've discussed here and highly conventionalized genres of the kind on which sociolinguists have focused. We have not tried to turn our theory of genre into a formal-pragmatic model that generalizes and expands on features of QUD models, but we think that this would be worth doing.[22]

Our theory also raises normative questions. How does a conversation's genre influence how it ought to be conducted—for example, by altering the norms governing speech acts? Who gets to choose a conversation's genre, and how are

[21] For related notions of pragmatic competence, see Unnsteinsson (2022) and Roberts (2022).

[22] If we were to do this, we would draw inspiration from models of related phenomena due to Bledin and Rawlins (2014) and Vesga and Starr (2025).

these decisions impacted by interlocutors' power differentials? Can a conversation's genre be unjust or oppressive? For example, if one participant can unilaterally impose a make-believe genre on the conversation when the other is trying to state important facts, could this amount to a pernicious kind of silencing? Again, it would take more space than we have to do justice to these questions, but we are optimistic that the framework that we've developed suggests a broad strategy for investigating them, since it suggests that they are special cases of broader questions about how norms and power dynamics interact with joint planning and collective action in general.

References

Armstrong, J. (2021). Communication before communicative intentions. *Noûs*, early view.

Bach, K. and Harnish, R. M. (1979). *Linguistic Communication and Speech Acts*. MIT Press, Cambridge, Mass.

Bade, N., Picat, L., Chung, W., and Mascaenhas, S. (2022). Alternatives and attention in language and reasoning:a reply to mascarenhas and picat (2019). *Semantics and Pragmatics*, 15(2):1–31.

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.

Biber, D. and Conrad, S. (2019). *Register, genre, and style*. Cambridge University Press.

Bledin, J. and Rawlins, K. (2014). What ifs. *Semantics and Pragmatics*, 12(14):1–62.

Boccaccini, M. T., Gordon, T., and Brodsky, S. L. (2005). Witness preparation training with real and simulated criminal defendants. *Behavioral sciences & the law*, 23(5):659–687.

Bratman, M. (1987). *Intention, plans, and practical reason*. Harvard University Press, Cambridge, Mass.

Bratman, M. (1992). Shared cooperative activity. *The Philosophical Review*, 101(2):327–341.

Bratman, M. (2014). *Shared Agency: A Planning Theory of Acting Together*. Oxford University Press, Oxford.

Broome, J. (2013). *Rationality Through Reasoning*. Wiley-Blackwell, West Sussex.

Buchanan, R. and Schiller, H. I. (2022). Pragmatic particularism. *Philosophy and Phenomenological Research*, 105(1):62–78.

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.

Cappelen, H. and Dever, J. (2019). *Bad Language*. Oxford University Press, New York, NY.

Carter, J. A. (2024). Knowledge and conversation. In Silva, W., editor, *Epistemology of Conversation*. Springer.

Carter, J. A. (2025). Knowledge norms and conversation. In Filho, W., editor, *Epistemology of Conversation: First Essays*. Springer, Cham.

Charlow, N. (2013). What we know and what to do. *Synthese*, 190(12):2291–2323.

Ciardelli, I., Groenendijk, J., and Roelofsen, F. (2009). Attention!'might'in inquisitive semantics. In *Semantics and linguistic theory*, volume 19, pages 91–108.

Clark, H. H. (1979). Responding to indirect speech acts. *Cognitive Psychology*, 11(4):430–477.

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

Clynes, A. and Henry, A. (2004). Introducing genre analysis using brunei malay wedding invitations. *Language Awareness*, 13(4):225–242.

Condoravdi, C. and Lauer, S. (2012). Imperatives: Meaning and illocutionary force. In Pinõn, C., editor, *Empirical Issues in Syntax and Semantics 9: Papers From the Colloque de Syntaxe et Sémantique à Paris 2011*, pages 37–58.

Coulthard, M. and Ashby, M. (1975). Talking with the doctor, 1. *Journal of Communication*.

Currie, G. (1986). Works of fiction and illocutionary acts. *Philosophy and Literature*, 10(2):304–308.

Currie, G. (1990). *The Nature of Fiction*. Cambridge University Press.

Dever, J. and Schiller, H. I. (2020). This paper might change your mind. *Noûs*, 55(4):863–890.

Ferguson, C. A. (1983). Sports announcer talk: Syntactic aspects of register variation. *Language in society*, 12(2):153–172.

Fisher, R. and Ury, W. (1981). *Getting to Yes: Negotiating Agreement without Giving in*. Houghton Mifflin.

Frankfurt, H. (1986). On bullshit. *Raritan Quarterly Review*, 6(2):81–100.

Friedman, D. C. (2025). Intentions and inquiry. *Mind*.

Friedman, J. (2013). Question-directed attitudes. *Philosophical Perspectives*, 27(1):145–174.

Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66(3):377–388.

Grice, H. P. (1969). Utterer's meaning and intention. *The Philosophical Review*, 78(2):147–177.

Grice, H. P. (1975). Logic and conversation. In Cole, P., editor, *Syntax and Semantics 3: Pragmatics*. Academic Press, New York.

Grice, P. (1989). *Studies in the Way of Words*. Harvard University Press, Cambridge, Mass.

Grosz, B. and Sidner, C. L. (1986). Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204.

Harris, D. W. (2014). *Speech Act Theoretic Semantics*. PhD Dissertation, City University of New York Graduate Center.

Harris, D. W. (2019). Intention and commitment in speech acts. *Thoeretical 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. (2021). Imperative inference and practical rartionality. *Philosophical Studies*.

Harris, D. W. (2025). Gricean communication, natural language, and human evolution. In Geurts, B. and Moore, R., editors, *Evolutionary Pragmatics*. Oxford University Press, Oxford.

Harris, D. W., Fogal, D., and Moss, M. (2018). 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.

Hawkins, R. D., AndreasStuhlmüller, Degen, J., and Goodman, N. D. (2015). Why do you ask? good questions provoke informative answers. In Noelle, D. C., Dale, R., Warlaumont, A. S., Yoshimi, J., Matlock, T., Jennings, C. D., and Maglio, P. P., editors, *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, Pasadena, CA. Cognitve Science Society.

Hoek, D. and Bradley, R. (2022). Million dollar questions: Why deliberation is more than information pooling. *Social Choice and Welfare*.

Hymes, D. (1977). *Foundations in Sociolinguistics*. Tavistock Publications.

Kaufmann, M. (2012). *Interpreting Imperatives*. Springer, Dordrecht.

Lewis, D. K. (1978). Truth in fiction. *American Philosophical Quarterly*, 15(1):37–46.

McGowan, M. K. (2019). *Just Words: On Speech and Hidden Harm*. Oxford University Press.

McKinney, R. (2016). Extracted speech. *Social Theory and Practice*, 42(2):258–284.

Moore, R. (2018). Gricean communication, language development, and animal minds. *Philosophy Compass*, 13(12):1–13.

Roberts, C. (1989). Modal subordination and pronominal anaphora in discourse. *Linguistics and Philosophy*, 12(6):683–721.

Roberts, C. (2012). Information structure in discourse: Toward an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5:1–69.

Roberts, C. (2018). Speech acts in discourse context. In Fogal, D., Harris, D., and Moss, M., editors, *New Work on Speech Acts*, Oxford. Oxford University Press.

Roberts, C. (2022). The architecture of interpretation: Dynamic pragmatics and pragmatic competence. *Unpublished manuscript of the 2022 Beth Lecture, University of Amsterdam*.

Schiffer, S. (1972). *Meaning*. Oxford University Press, Oxford.

Searle, J. (1975). The logical status of fictional discourse. *New Literary History*, 6(2):319–332.

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.

Stokke, A. (2023). Fictional force. *Philosophical Studies*, 180:3099–3120.

Strawson, P. F. (1964). Intention and convention in speech acts. *The Philosophical Review*, 73(4):439–460.

Thomasson, A. L. (1998). *Fiction and Metaphysics*. Cambridge University Press, New York.

Unnsteinsson, E. (2017). A gricean theory of malaprops. *Mind and Language*, 32(4):446–462.

Unnsteinsson, E. (2022). *Talking About: An Intentionalist Theory of Reference*. Oxford University Press, Oxford.

Unnsteinsson, E. (2023). Authentic speech and insincerity. *Journal of Philosophy*, 120(10):550–576.

van Rooy, R. (2003). Questioning to resolve decision problems. *Linguistics and Philosophy*, 26:727–763.

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

Vesga, A. and Starr, W. (2025). Non-literal communication and practical coherence. In Preyer, G., editor, *Coherence in Discourse*. Oxford University Press, Oxford.

Williamson, T. (1996). Knowing and asserting. *Philosophical Review*, 105:489–523.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.

Worsnip, A. (2021). *Fitting Things Together*. Oxford University Press, Oxford.

Yalcin, S. (2007). Epistemic modals. *Mind*, 116(464):983–1026.

Yalcin, S. (2018). Expressivism by force. In Fogal, D., Harris, D., and Moss, M., editors, *New Work on Speech Acts*. Oxford University Press.