Genre and Conversation

Daniel W. Harris and Elmar Unnsteinsson

Draft of October 9, 2023

There are different genres of conversation. Joking around with friends is different from a serious conversation with high stakes, and both are different from a tense, adversarial deposition. A conversation's genre affects what it is natural to say, and how it is natural to interpret what others say. In a practical, decision-making conversation, it is difficult to merely state a fact without implying that it is relevant to the decision at hand. If we are engaged in a brainstorming session, it is permissible to say things that we aren't sure about, and that may not be compatible with what others have said. In a game of make-believe, it may be difficult to convey information about the real world.

In this paper, we propose to think of genres as recurring patterns in conversation plans, which are the complex 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 exchange 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 goals, and we are under rational pressure to do so in a coherent 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, giving balance to conversations.

1 What are Genres of Conversation?

Conversational genres have mostly been discussed within sociolinguistics. An influential and early characterization comes from Dell Hymes: "genres…are categories such as poem, myth, tale, proverb, riddle, curse, prayer, oration, lecture,

commercial, form letter, editorial, etc" (1977, 61). Sociolinguistics has 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. In addition, they tend to focus on highly conventionalized genres, where participants may be following strict cultural scripts with specified roles and rituals.

We think that these conventional genres are real and interesting. But we want to focus on what we think is a less conventionalized kind of genre variation in this paper. In particular, we'll focus on the following four axes along which conversations can differ with respect to genre:

(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 saying that these genre distinctions are non-conventional, we don't mean that there aren't culturally specific conventions about how to conduct oneself within each genre. For example, any given population of speakers might have its own linguistic conventions for how to signal that it is engaging in make-believe, or seeking to be cooperative. But theses are conventions about the implementation details rather than the nature of the genres themselves. Although every speech community may have its own unique way of indicating that they are engaging in a practical conversation, and perhaps its own grammatical mechanisms for doing so, we assume that all human speech communities sometimes engage in practical conversations and sometimes in informative ones. We can say the same about the other genres. This stands in contrast to sports-announcer talk or doctor-patient interviews, whose very existence is owed to culturally local conventions.

Philosophers and linguists have tended to focus mostly on conversations which are best described by the first item of each contrasting pair.[1] So, the norm is to focus on informative, committal, factual, and cooperative exchanges. For example, although we think that Grice's (1975) cooperative principle is operative in conversations of all of these types (perhaps with the exception of extremely adversarial conversations), his maxims of conversation are formulated in such a way that they only neatly apply to informative, committal, factual, and cooperative conversations. Grice was aware of this, and said that we would need different versions of the maxims for other kinds of conversations:

> 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)

Grice didn't follow up on this point, and surprisingly little has been said about it by others. We will try to say more in what follows.

We'll spend the rest of this paper answering several questions about genre that have been widely neglected. In the next four subsections, we will take a closer look at our four non-conventional genre distinctions, and try to get a bit more of a pre-theoretic grip on them. What all of these genres share, we'll argue, is that they make certain kinds of speech acts more natural and expected than others. This will be the explanandum, whose explanans will come in later sections, where we will tackle some theoretical questions: What are conversational genres? How do they arise? Why do they exist?

1.1 Informative vs. practical

In informative conversations, we share information, whereas in practical conversations, we make decisions about what to do. Directives, suggestions, proposals,

[1] 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 conversations as joint activities, which is similar in explanatory orientation but adds variations which we think have to do with local implementation. 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).

and questions are the default speech acts in practical conversations, and assertions are fitting or natural only insofar as they bear on the decision at hand. Similarly, assertions and questions about matters of fact are likely to be interpreted as indirect directives, suggestions, or proposals to act. Suppose, for example, that we are negotiating about how to divide our parental chores this evening. You see that I have already started helping the kids go to bed, and I say, "The dishwasher is full." This would normally be intended and interpreted as a contribution to our decision-making process, and not merely to share information. We will argue that genre is part of the explanation of why the speech act is indirect in this case.

To take another example, imagine that we are trying to figure out where to go for dinner and you say (5).

(5) Calexico has good tacos.

If (5) belongs to a practical conversation, it will be expected to have some practical or directive significance. It would be unnatural for you to follow up by saying that we should definitely not eat at Calexico, unless you have suddenly changed your mind. And this is part of what we want to explain, namely, how genre seems to make it more natural to understand a simple assertive utterance as an indirect suggestion or proposal to do something.

1.2 Committal vs. exploratory

In committal conversations, utterances are treated as firm commitments. It is often natural to ask the speaker how they know, or what their evidence is, when they have said something. It can be equally fitting to criticize a speaker for speaking falsely. By contrast, in exploratory conversations, we merely make suggestions, float possibilities, and try out ideas, without commitment. Exploration allows that speakers can say things which directly contradict what others have said, or even what they themselves have previously said. Asking for justification or expressing disagreement tends, however, to be infelicitous or ill-fitting.

The practically-oriented literature on negotiation contains many insights into exploratory genres of conversation. One of the most influential and widely used textbooks in this area is called *Getting to Yes* (Fisher and Ury 1981). The authors divide negotiations into steps, one of which is a brainstorming session,

> [...] 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.
>
> In a brainstorming session, people need not fear looking foolish since wild ideas are explicitly encouraged. And in the absence of the other side, negotiators need not worry about disclosing confidential information or having a half-baked idea taken as a serious commitment. (Fisher and Ury 1981)

Now, strictly speaking, it is not true that people do not pause to consider whether something is unrealistic. We can easily imagine a speaker who only ever proposes manifestly wild or impossible ideas, making them useless for brainstorming. The truth, however, is that exploration is all about filling the space of possibilities, by proposing something that *might* be done or *might* be true.

To take another example, consider a group of doctors brainstorming about a difficult diagnosis, writing their proposals on a whiteboard. Scenes of this kind are common in the TV series Dr. House. One of the doctors might say:

(6) It's lupus.

And then another might immediately offer a different, inconsistent diagnosis. The whiteboard might even end up containing a direct contradiction, "It's not lupus." Here it seems like the utterance of (6) has an implicitly modal flavour. The genre makes it natural to understand the utterance as suggesting a possibility. A signature of exploratory conversations is the fact that contradicting what has been said before is relatively natural. True, no one in the group thereby endorses anything of the form *p and not-p*, not even as a possibility. But in a committal conversation, inconsistent utterances are treated very differently.

1.3 Factual vs. make-believe

In factual conversations, our aim is to say true things about the actual world. But when we engage in make-believe—telling stories or accepting something for the sake of the conversation, for example—the relevant norms of truth and knowledge are not in place. When the speaker makes a supposition or says something which is literally false—"There was once a woman who had a third eye in her bellybutton, her name was Gilitrutt"—it does not fit the genre to ask, "How do

you know?" or "That's false." Suppose that I am telling you a made-up story and I utter:

(7) Gilitrutt has three eyes.

In this case, I want you to join me in taking the attitude of pretense, rather than belief, toward the content of (7). This is indicated, for example, by the way in which you would be missing the point if you replied, "I don't believe you!"

Accordingly, genres of make-believe have distinctive effects on default interpretation. Roughly, make-believe tends to suspend real-world knowledge, which we employ automatically in interpretation and, often, the pretense will contain gaps. Let us stay with a the simple example of storytelling. Within the story about three-eyed Gilitrutt, I might say, "Then she glanced momentarily at her watch, sighing with obvious irritation." Normally, we would assume that glancing is performed by the two eyes in the head, but Gilitrutt might very well be using her bellybutton eye, and the audience might be intended to recognize this from contextual cues. Crucially, this is not what happens when we are just describing highly unusual situations because, in that case, interpretation can in principle rely on real-world knowledge. Now, it is controversial whether all fiction is a form of make-believe, and our theory is silent on that issue, because make-believe is a much broader phenomenon which occurs in a range of different settings.[2] To accept something for the sake of the conversation, in the Stalnakerian sense, is a form of make-believe too, where norms of truth and knowledge are not in force (e.g., Stalnaker 1984).

1.4 Cooperative vs. adversarial

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, because we can't count on our interlocutors to interpret us charitably. So it is risky to use ambiguous expressions or to communicate indirectly, as an addressee will be more likely to attribute an unintended meaning. Suppose, for example, that Steven and Anna are having an acrimonious conversation about whether Steven has recently been changing his share of their child's diapers.

[2] Although, for example, Currie (1990) argues for a make-believe theory of fiction, Lewis (1978) and Thomasson (1998) would say that fiction involves ordinary belief and that the differences are due to the nature of the contents believed.

(8) Anna: It just doesn't seem to me that you're doing your share.
Steven: I have just been busy with really important things lately at work.
Anna: Are you implying that I haven't!?
Steven: No!

Let's suppose that Steven didn't, in fact, intend to convey that his recent tasks have been more important than Anna's. Maybe he merely intended to convey that he was struggling to get used to his new duties at work, and offering this as an explanation for his slacking off at home. In a less acrimonious context, Anna might have attributed something like this indirect meaning to Steven. But in the context of a relationship fight, charity can be hard to come by, and so it was a bad idea for Steven to leave so much of his message implicit.

In some adversarial contexts—typically, those in which everyone knows about the sources and nature of the adversariality—we are less likely to interpret speakers as having implicated anything at all. 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 were up to yesterday—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 think that this might be relevant to A, they might 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, the fact that they didn't volunteer this information suggests that they 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, speculated that B must have had some ulterior motive for withholding information, and concluded 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. In this sort of context, nobody is expected to volunteer more information than they have to,

and everyone knows that it would be foolish for B to respond as in (10). 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.

2 Planning Theory

We're going to argue that conversation genres are recurring patterns in conversation plans. So we'll need a theory of what a conversation plan is. We'll build this theory on top of Bratman's (1987; 2014) planning theory.[3] Because Bratman's theory of shared plans is built as an extension of his theory of how individuals make plans, it will be easiest to explain that first, before moving on to the shared case.

On Bratman's view, plans are hierarchical structures whose elements are intentions. 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 that goal. In order to work out these implementation details, the agent must engage in practical reasoning in order to 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 should take place, who to invite, and what to serve. Suppose that they decide to serve pasta. Even this subplan leaves many questions unanswered: Where will they get the ingredients, how exactly 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 goal with which they began to specific motor instructions that can initiate specific bodily movements.

An agent who is planning well will be responsive to rational pressure to form new intentions in a way that is consistent with their other intentions, with their beliefs about what they can do, and with the need to flesh out the details of 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.

This sort of planning carries significant advantages for creatures like us, who

[3]We're working with Bratman's theory because it is the most influential theory that explains joint action in terms of plans, but also because it works very well for our purposes. It is an interesting question whether and how rival theories of joint action could be plugged into the rest of our story, but we won't try to pursue that question.

are cognitively limited but who aspire to complex actions. It allows us to break complex, multivariate decisions into tractable chunks. Because the outcomes of earlier decisions constrain the set of options from which we must 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 we otherwise could with the limited cognitive resources that we have available, and to pursue goals that are more abstract, in the sense that they would have to be pursued by different specific actions in different circumstances.

Bratman's theory of shared agency uses ingredients from his planning theory to develop an austere account of what it is to make shared plans and act together. The main new ingredient is for two agents to have a shared intention, which on Bratman's view requires that each intends some end, intends to pursue this end by means of meshing subplans, and is aware that the other has these intentions. To have meshing subplans of a shared intention is to have intersubjectively coherent intentions about how to accomplish it that together add up to a complete representation of a way of pursuing the intention. For example, if Dan and Margot 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 Margot forms an intention to plan the event and an expectation that Dan will prepare the meal, then Dan should form the converse intention and expectation.

A more detailed example of this kind is illustrated in Figure 1. Consider, briefly, several features of this planning structure that are common to many shared plans. First, notice that the shared intentions have unshared intentions both as subplans and as superplans. They tend to have unshared intentions as subplans because it is ultimately individuals who must translate the shared plans into particular bodily movements, such as picking up the telephone to call Justin. Thus Bratman's requirements about meshing subplans of shared intentions, which are typically not themselves shared. On the other hand, shared intentions typically have unshared intentions as superplans because individual agents typically have their own, private reasons for engaging in joint activities. Dan's reasons for intending to have a dinner party needn't be the same as Margot's for example.

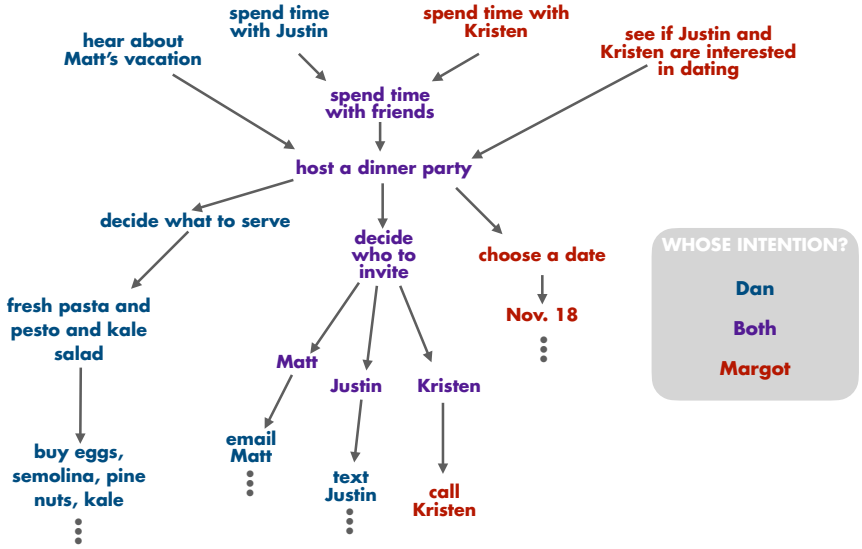

Figure 1: An example joint plan, with arrows pointing from intentions to their subplans.

3 Conversation Plans

Conversations are themselves joint activities that are governed by shared plans.[4] These are what we call "conversation plans." Figure 2 is a schematic illustration of a conversation plan that includes some of the elements that are common (though not universal) to the planning structures governing conversations, and that will be important in the sections that follow.

At the top of the diagram are the individual intentions that interlocutors are seeking to pursue by means of the conversation. For example, suppose that Dan and Elmar are having a conversation about where to eat dinner. Dan intends to eat a low-calorie meal, while Elmar intends to carb-load for the half marathon that he is running tomorrow. These intentions aren't shared, but because they represent the goals that Dan and Elmar are trying to achieve by having a conversation, they will figure into how the conversation plays out, because further elements in the conversation plan will have to cohere with and further them. Next there are the various shared intentions that govern our conversation, which interlocutors form as a joint means of pursuing their individual ends. For example, Dan and

[4]For some influential earlier discussions of conversation as a jointly planned activity, see Bratman (2014); Carlson (1982); Clark (1996); Cohen and Levesque (1985, 1990); Ginzburg (1995a,b, 2012); Grice (1975); Grosz (1986); Perrault (1990); Roberts (2012b); Thomason (1990).

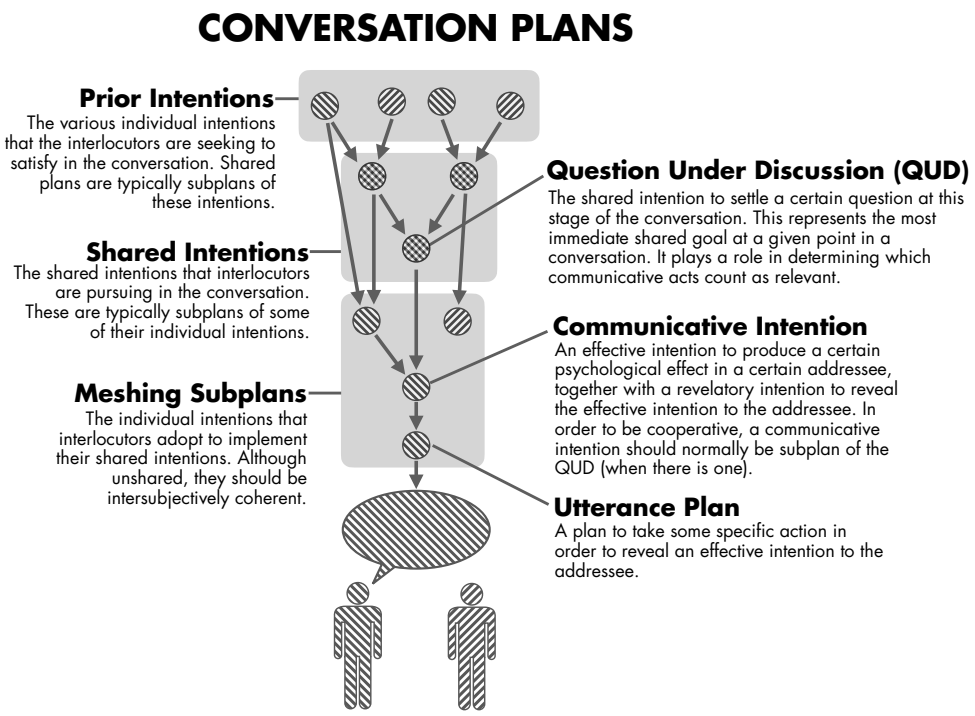

Figure 2: A schematic conversation plan, highlighting some of the kinds of elements that often (though not always!) show up in conversation plans.

Elmar each intend that they will eat together, and that they will work out how to do this by means of a polite conversation, and they intend to pursue these ends by means of meshing subplans. These elements in the conversation plan meet the definition of shared intentions given above. These shared intentions will often include an intention to answer what Roberts (2012b) calls the "immediate question under discussion," or "QUD." Answering this question is the interlocutors' most immediate shared conversational goal, and an utterance will normally be considered relevant only if it helps to answer or refine the QUD (more on this in §4. Below the level of shared plans, there are interlocutors' meshing subplans, which are their individual but (ideally) coordinated intentions about how to accomplish their shared goals. These may include things like a listener's intention to give the speaker a chance to say something, the intention to pay careful attention to what the other is saying, and so on. Most importantly, there are communicative intentions, which determine what speakers mean, and so determine the success conditions of individual communicative acts. For example, if Elmar

says, "it's closed" with a communicative intention to inform Dan that Calexico is closed, then he will have performed an assertion.

In the next several sections, we will argue that conversation plans really do include elements of these kinds, and that conversational genres are best understood as recurring patterns in the structure of conversation plans.

4 Questions Under Discussion as Elements in Conversation Plans

We begin our discussion of conversation plans in the middle, with agents' shared plans to answer QUDs. Why think that we form such plans? Our answer, in short, is that some of our most robust empirical predictions about how conversations work have come from QUD-based pragmatic models, and we can make the most sense of these models if we take their subject matter to be communicators' shared plans.

QUD models have helped to generate useful predictions about a wide range of semantic and pragmatic phenomena.[5] Most importantly, consider exchanges like (11) and (12):

(11) (a) Elmar: Where should we eat?
(b) Dan: We should eat at a Mexican restaurant.
(c) Dan: # We should try scuba diving.

(12) (a) Elmar: Where should we eat?
(b) Dan: Should we try the new dumpling place?
(c) Dan: # What was Wes Anderson's first movie?

In each of of these exchanges, Dan responds relevantly to Elmar's question with (b), but not with (c). Intuitively, this is because Dan's (b)-responses address the question, either asserting an answer (as in (11)) or asking a further question that refines Elmar's question in a useful way (as in (12)).

With her QUD model, Roberts (2012b) devised a way to both generate precise and widely accurate predictions about relevance judgments of this kind, and to

[5] See, *inter alia*, work on projective content (Simons et al., 2017, 2010), loose talk and metaphor (Hoek, 2018), disjunction (Simons, 2001), epistemic modals (Beddor and Egan, 2018), attitude verbs (Schaffer, 2007; Yalcin, 2018), and our ability to interpret semantically underspecified expressions in general (Schoubye and Stokke, 2016).

explain something about why we have them when we do. She models conversations as revolving around shared contexts that include both a *common ground* and a *question stack*. Following Stalnaker (1978; 2002), Roberts thinks of the common ground at a given stage of a conversation as the set of propositions that the participants are taking for granted for the purposes of the conversation at that stage. This allows us to likewise represent the context set, which is the set of possible worlds at which all of the propositions in the commmon ground are true. Stalnaker suggested that many conversations could be thought of as what Roberts calls "joint inquiry" (Roberts, 2012b, 4), in which a collection of two or more interlocutors work toward a shared goal of pooling their information by making assertions. This process can be modeled as the addition of propositions to the common ground, thereby whittling away candidates for actuality from the context set.

We are normally pursuing conversational goals more specific than asserting any new information whatever, and this is where the question stack comes in. This stack is an ordered list of question-contents, each of which is modeled as a set of propositions—intuitively, the set that includes all of the question's mutually exclusive complete answers.[6] The question at the top of the stack is the immediate question under discussion, or QUD. When someone asks a question in a conversation (and nobody objects), the content of their question is installed at the top of the stack, becoming the new QUD and pushing the old question down a spot. During its tenure as QUD, a question partitions the context set into a set of contextually live possible answers, inviting the interlocutors to assert something that eliminates some of them—thereby partially answering the QUD—or that is incompatible with all but one—thereby completely answering the QUD. When someone completely answers the QUD in this way, it is eliminated from the stack, promoting the next unanswered question down to be the new QUD. Roberts thus predicts that an assertion is relevant if and only if it at least partially answers the QUD. On the other hand, she predicts that a question q is relevant if and only if its content is a subquestion of the QUD, which is to say that a complete answer to q would be a partial answer to the QUD. This lends precision to the idea that a relevant conversational move is one that makes progress on the QUD, either by answering it or by pursuing a "strategy of inquiry" that, if successful, will help us to answer it.

[6]This way of modeling question-contents is motivated by the dominant tradition of work on the semantics of interrogative clauses (Groenendijk and Stokhof, 1984; Hamblin, 1973; von Stechow, 1991).

It is important to note that it is sometimes felicitous to make a communicate act that isn't *literally* relevant to the QUD, so long as one indirectly communicates something that is relevant. Consider the following example:

(13) (a) Elmar: Where should we eat?
(b) Dan: Calexico has good tacos.

In this case, Dan's response probably doesn't meet Roberts' definition of relevance.[7] But in this case, it is tempting to attribute to Dan not just the communicative intention to inform Elmar that Calexico has good tacos, but also an additional communicative intention to get Elmar to consider going there. Roberts' model actually helps to explain why we are tempted to attribute this indirect speech act to Dan. If Dan had merely asserted that Calexico has good tacos, then his contribution would not have been relevant. But if Elmar assumes that Dan is a rational, cooperative interlocutor, then he will conclude that Dan must not *really* have been speaking irrelevantly. He must have meant something additional, which *was* relevant. Since a restaurant's having good tacos is a reason to eat there, and since the QUD is *Where should we eat?*, the most obvious candidate for Dan's indirect speech act in this situation was a suggestion that they should eat at Calexico. As Roberts acknowledges, this is very similar to Grice's (1975) explanation of how we derive relevance implicatures. But since Roberts gives precise content to the concept of relevance, her theory comes closer to making specific predictions about which kinds of implicatures we should expect in particular cases.

Roberts' model also generates impressively accurate predictions about the appropriateness conditions of prosodic focus, which is the way in which we stress certain words or syllables of our utterances to a greater degree than others. Building on work by von Stechow (1991), Roberts shows that which sorts of prosodic focus are felicitous is itself a function of the QUD.

(14) S: Who invited Sue?
A1: $[\text{Mary}]_{\text{F}}$ invited her.
A2: # Mary invited $[\text{her}]_{\text{F}}$.

[7] I say "probably" because Dan's response would be literally relevant in a context in which it was common ground that there is only one restaurant with good tacos and they will eat at whichever restaurant that is. In this case, the content of Dan's assertion would contextually entail an answer to the QUD.

(15) S: Who invited whom?
A1: Who did [Mary]$_F$ invite?
A2: # Who did Mary [invite]$_F$?

In exchanges like these, S's question establishes a new QUD. Roughly speaking, responses to the question that include prosodic focus are felicitous if and only if they would have the QUD as their content if the focused element (here marked with $[\cdot]_F$) were replaced by a *wh*-phrase. Even more roughly, this is to say that an answer is congruent with a question if and only if the focused element in the answer corresponds to the *wh*-phrase in the interrogative clause used to ask the question.[8] These turn out to be very robust generalizations about when it is appropriate to use prosodic focus. Roberts explains this phenomenon by arguing that prosodic focus is a grammatical mechanism whose function is to test and reinforce interlocutors' coordination on the QUD.

We'll take it for granted that QUD-based models of conversation are successful, and that this success calls out for an explanation in terms of an independently motivated specification of the model's worldly subject matter. In other words, what features of real conversations explains the fact that Roberts' model makes such accurate predictions? Our answer, which fleshes out Roberts' own remarks, is that whereas the common ground is an idealized model of the beliefs and other belief-like states that interlocutors rely on when designing and interpreting speech, the question stack is an idealized model of our plans to resolve questions.[9]

An alternative interpretation of QUD models would be to think of them as

[8]For the more technically precise version of these ideas, see Roberts (2012b, §2.1).

[9]Roberts introduces her theory by telling us that "[d]iscourse is organized around a series of conversational goals and the plans or strategies which conversational participants develop to achieve them (Roberts, 2012b, 3). She also sometimes refers to QUDs as "discourse goals" (Roberts, 2012b, 26), distinguishing these from "domain goals," which are "things we want to achieve quite apart from inquiry" (Roberts, 2012b, 7). In later refinements of her model, Roberts describes discourse goals as a "distinguished type of domain goal, those the interlocutors are jointly committed to achieving in the discourse itself" (Roberts, 2018, 323). In her original paper positing the question stack, Roberts does not directly cite Bratman as an influence on her theory, but she does repeatedly mention "Planning Theorists in artificial intelligence," citing a number of authors who themselves drew on Bratman's theory of intentions and practical reasoning in building computational models of discourse—e.g. Cohen and Levesque (1990); Cohen and Perrault (1979); Grosz (1986); Perrault (1990); Thomason (1990). In the Afterward to her paper, which was written years later, Roberts explains that she was influenced by Bratman's work while she was a postdoc at Stanford in 1986–1988, where he was developing his theory at the time (Roberts, 2012a, 3).

models of something essentially *grammatical*—of some computational machinery internal to the faculty of language, say.[10] But there are good reasons to doubt this interpretation. Roberts argues that there are nonlinguistic routes by which a question can become the QUD. Just as some fact can become common ground because we all observe each other observing it, a question can become the QUD if we all sense each other wanting to figure out the answer. If a goat walks into the room, it will tend to become common ground that there is a goat in the room (Stalnaker, 1999, 86), but it is also likely that their conversation will turn to the question, "Why is there a goat in the room?", even if nobody explicitly asks it. This suggests that QUDs are models of some not-essentially-linguistic aspect of our psychology, even if we do have specialized linguistic means for manipulating them.

So, why should we think of QUDs as modeling shared plans, specifically? The first and most important reason is that this interpretation gives an independently motivated explanation of the relationship between QUDs and relevance. 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 answer the QUD. Suppose that Elmar and Dan have a shared intention to answer the question, *Do they have horchata at Tacombi?* In this context, if Dan knows the answer, it makes sense for him to pursue their common goal by forming a communicative intention to add this information to the common ground. The fact that this communicative intention is a constructive and coherent subplan of their intention to answer the QUD is what makes Dan's assertion relevant. By contrast, if he were to form a communicative intention to assert that Wes Anderson's first film was Bottle Rocket, his intention would *not* be a constructive subplan of their shared plan. This explains our judgment that *this* assertion would be *ir*relevant.

We can explain relevance judgments about questions in a similar way. Suppose that Elmar and Dan have a shared plan to answer the question, *where should we eat?* In that context, Elmar might decide to pursue this shared plan by forming a communicative intention whose aim is to establish a new shared intention to answer the question, *should we eat at Calexico?* This speech act will be relevant because the plan that it proposes is a good means to Elmar and Dan's prior end of answering their prior QUD. In general, posing a subquestion of the QUD will often be relevant because answering a subquestion is a good means to the end

[10]Several authors have recently argued that we should think of formal models of context in terms like this (Lepore and Stone, 2015; Stojnić, 2018).

of making progress on the question. This explains why contiguous questions in Roberts question stack are supposed to stand in question-subquestion relations.

On the interpretation of the QUD model that I have just proposed, Roberts' notion of relevance turns out to be just a special case of the broader phenomenon of plan coherence. To make a relevant move in a conversation is just to do something with a communicative intention that is a constructive and coherent subplan of the interlocutors' shared plans. There is nothing particularly unique about conversation, on this view. It is, as Grice foretold, just one kind of cooperative, rational activity.

5 Genres as QUDs

We are now in a position to understand the first two genre distinctions that we discussed in §1. For conversations to vary along these dimensions, we think, is just for them to vary in the kind of QUDs that their participants are trying to answer.

First consider our distinction between informative and practical conversations. Intuitively, this is a distinction between conversations in which interlocutors are trying to share information and conversations in which they are trying to decide what to do. We operationalized this distinction by pointing out that these different kinds of conversation make different kinds of speech acts more or less natural or expected. Practical conversations are the natural habitat of directives, suggestions, proposals, and questions about what to do. In the context of this sort of conversation, an assertion will be welcome only insofar as it bears on the practical matter at issue, and assertions will often be understood as indirect directives or proposals. By contrast, in an informative conversation, utterances that would literally be interpreted as directives will tend to be understood as indirect assertions.

We can easily explain this data in terms of differences in QUD. Consider again this example that we discussed above:

(5) (a) Elmar: Where should we eat?
    (b) Dan: Calexico has good tacos.

Here, Dan utters a declarative sentence—a sentence whose literal use would be to perform an assertion—but it is natural to interpret him as making an indirect proposal about what to do. The obvious explanation of this is that by Elmar asking a practical question—a question about what to do—Elmar turned their

conversation into a practical one. Practical conversations, on this view, are just conversations with practical QUDs.

By contrast, an informative conversation is one whose QUD is an informative QUD—a question about what to believe. To illustrate how a QUD can render a conversation informative, consider the following exchange:

(16) (a) Elmar: How do people get to Grand Central Station?
(b) Dan: Take the 5 train.

Here Dan utters an imperative sentence—one that would normally be used to perform a directive act. But it is clear in this context that Dan's point is merely to give Elmar information about how people get to Grand Central, and not to try to get Elmar himself to take the 5 train. This is what some authors have called an instructional or disinterested-advice use of an imperative—one whose point is informative rather than directive (e.g. Kaufmann 2012, 141). These kinds of utterances are characteristic of informative conversations. But now we can say that what it is for a conversation to be informative is just for it to have an informative QUD.

We can give a similar account of the distinction between committal and exploratory conversations. Specifically, our claim is that committal conversations are those in which the QUD is a question about how things are or what interlocutors want to do, whereas exploratory conversations are those in which the QUD concerns how things *might* be or what interlocutors *might* want to do. Consider the following brainstorming exchange:

(17) House: What might be causing this patient's symptoms?
Thirteen: A brain tumor is restricting blood flow to her hypothalamus.
Kutner: She fell and hit her head.
Taub: She has lupus.
Thirteen: She was exposed to a toxin.

Given the context, these utterances won't naturally be understood as assertions—i.e., as attempts to get addressees to believe their contents, or to add their contents to the common ground. The point of these utterances, we submit, is merely to say what might be the case—to populate the space of possibilities, so that the group can go on, in a later stage of the conversation, to reason about which of these options best explain the patient's symptoms. We can see this from the fact that there is nothing incoherent about Thirteen suggesting two incompatible diagnoses, for

example. After all, the current conversational goal is merely to get some live options on the table before deciding which one is right. In general, we can induce this sort of context by asking a might-question, as House does in (17). But of course, the doctors on *House* have a conversation of this kind in every episode, and so they know that this is the kind of conversation they're having—i.e., that this is the kind of QUD they're answering—even if House doesn't explicitly pose an epistemically modalized question.

Given our argument in §4 that QUDs are models of shared plans, the upshot of this section is that at least some genre differences should be understood as different kinds of shared plans. If *all* genre distinctions were just QUD distinctions, our claims about conversation plans would just amount to the claim that QUD models can explain yet another phenomenon, together with an interpretation of those models. But we don't think things are quite so simple.

6 Genres as Shared Superplans of QUDs

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

(18) Dan: How many eyes does Gilitrutt have?
Elmar: Three.

This exchange could belong to two different conversational genres. On one hand, it might belong to a game of make believe, in which case both Dan and Elmar understand that they are making things up that aren't actually true. Even if Elmar's utterance is wholly successful, Dan won't come away believing that Gilitrutt exists, or has three eyes. Rather, the success of Elmar's utterance would consist in getting Dan to imagine, or make believe that Gilitrutt has three eyes. In this scenario, it would be confused for Dan to respond to Elmar by pointing out that what he says isn't true. That misses the whole point of the conversation.

On the other hand, the very same exchange could take place in a factual conversation, if only Dan and Elmar are both under the misapprehension that Gilitrutt is a real person, and that Elmar is imparting factual information to Dan. In this case, it would be perfectly natural for Dan to ask Elmar how he knows, or to express doubt about the accuracy of Elmar's claim. And in this case, Elmar's claim will be fully successful only if Dan winds up believing what Elmar has said.

Notice, however, that the QUD is the same in both cases, imposing the same relevance constraints on subsequent conversational moves. This suggests that

our distinction between factual and make-believe conversations should not be understood in terms of differences in the QUD itself. How, then, should we explain it?

Our proposal is that we should understand this genre difference in terms of Elmar's and Dan's reasons for adopting their QUD in the first place. In the factual case, they are seeking to answer the QUD as a way of gaining information about the world. In the make-believe case, they are seeking to answer the very same QUD, but for other reasons. For example, perhaps they are seeking to entertain themselves, or to cook up vivid imagery. Whatever it is that they are seeking to accomplish by engaging in make believe, this is what explains their plan to address the QUD. So whereas the genre distinctions that we addressed in §5 could be understood as different kinds of QUDs, the distinction between factual and make-believe conversations must be drawn above the level of the QUD, in terms of the more abstract plans in virtue of which interlocutors decide to address certain questions in the first place. The QUD is a subplan of these plans. Conversely, they are superplans of the QUD.

This theory of the factual/make-believe distinction probably generalizes to a range of ways in which conversations differ, which correspond to our various shared reasons addressing questions. Sometimes we want knowledge. Other times, we merely want to speculate, or bullshit with friends, or imagine things together. In these cases, we ask and answer the same questions, but the cognitive results of this will be different in different cases. We leave some conversations believing what was said, while other conversations may leave us only with new things to imagine. It needn't follow that the latter conversations are any less successful than the former. The two kinds of conversation were organized around different goals, which they may have pursued equally successfully.

We see this line of thoutht as closely related to Yalcin's discussion of what he calls "conversational tone," which he explains as follows:

> Given only what is common ground among a group of agents, one does not yet know how the agents of the context mutually regard the propositions in the common ground with respect to their other cognitive attitudes. To be given the common ground is only to be given a set of propositions mutually understood to be presupposed; it is not yet to be given that the agents also regard those presuppositions as knowledge, or as warranted belief, or conjecture, or fiction, or whatever. Using the notion of common ground, we can define a second notion which will let us articulate the status that the agents

of a given context attach to the propositions they presuppose. Call this notion *conversational tone*:

> An attitude is the *conversational tone* of a group of interlocutors just in case it is common knowledge in the group that everyone is to strike this attitude towards the propositions which are common ground.

(It may be that a conversation is plausibly understood as having more than one conversational tone, but let me focus on the case where there is just one. And let me stipulatively exclude presupposition itself from the class of possible conversational tones.) When interlocutors coordinate on a conversational tone, they come into agreement about what counts as the correct non-public attitude to take towards what is common ground. This will be a reflection, *inter alia*, of the purpose of the discourse. If the conversational tone of our discourse is knowledge, then we regard our common ground as common knowledge, and we take our discourse to be trafficking, and aiming to traffic, in factual information. Similarly with belief. If the conversational tone is pretense, then we are not attempting to keep the common ground compatible with the truth, and we take ourselves to be trafficking in fiction. And so on, for all the various attitudes around and in between: the conversational tone may be belief, or suspicion, or supposition, or high- credence-that, or ironic non-belief, etc., depending on the interests and purposes of the interlocutors. It may also be a conditional attitude: the conversational tone may, for instance, be belief (in each q in the common ground) *conditional* on some specified p. (Yalcin, 2007, 1008)

In positing common ground, Stalnaker created a piece of pragmatic machinery that turns out to work in roughly the same way in a surprising range of conversations. Common ground is a body of information, but it may be treated by interlocutors as what they commonly know, or as what they merely accept for the purposes of a pretense. Either way, Stalnaker's model makes similar predictions about which conversational moves are licit, and these predictions turn out quite well, in ways that are relatively insensitive to conversational tone.

In effect, we have been making a similar point about the QUD, which generates the same, mostly correct predictions about the (ir)relevance of utterances,

whether interlocutors are trying to answer it in the service of a serious conversation, a bull session, or a game of make believe. We think that both this point about the QUD and Yalcin's claims about conversational tone are best understood in terms of conversation plans that include shared intentions above the level of the QUD—shared intentions to do things like exchange knowledge or engage in make believe.

7 Genres as Patterns of Coherence and Conflict

Finally, let's consider the distinction between cooperative and adversarial conversations. Cooperativity 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 background information. Our claim here is only that one very important component of cooperativity as it shows up in conversation involves the degrees of both coherence and transparency in interlocutors' intentions. Specifically, we have in mind the individual, or unshared intentions that that they are seeking to accomplish by having conversations with each other. Insofar as these intentions are mutually supportive and transparently known to all, the ensuing conversation will tend toward cooperativity. To the extent that interlocutors' intentions conflict in relevant ways, the conversation will tend to be adversarial. And to the extent that interlocutors simply aren't aware of each other's goals, their ability to cooperate will be hamstrung.

Of course, each agent may intend to achieve a whole range of goals in a conversation, and any of these may be compatible or incompatible with the intentions of their interlocutors. Some instances of incoherence may turn out to be more relevant than others to the conversation. The quantity and relevance of incoherence in interlocutors' intentions will both contribute to the degree to which a conversation seems antagonistic, and so the cooperativity/antagonism distinction turns out to be a scale with more than one interacting dimension. And the degree to which a conversation seems antagonistic will also increase with the degree to which conflicts and their relevance is transparent.

We can say all of these things about joint activities in general, and not just conversations. Consider again the example of Margot and Dan planning a dinner party, which we discussed in §2. Suppose that Margot is planning the dinner party because she intends to introduce two friends whom she thinks would make a good couple, but Dan thinks that this would be a terrible idea, and intends to keep these two friends from ever meeting. This conflict in Margot's and Dan's

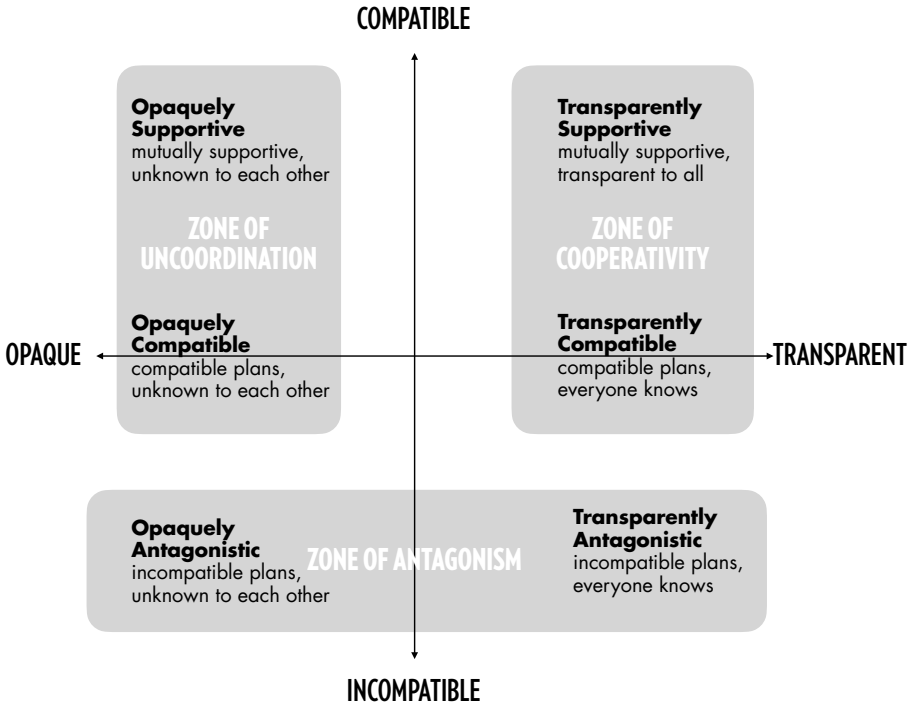

Figure 3: Two dimensions along which the private superplans of a shared intention can vary. Two agents' intentions can range from being mutually supportive to compatible to incompatible, on one hand, and from transparent (in the sense of being mutually known) to opaque, on the other.

private intentions isn't relevant to many aspects of the dinner-party-planning process. For example, it might not affect what they choose to serve. But when it comes time to decide whom to invite, it will likely lead to an impasse, as Margot will likely suggest inviting the two friends and Dan will demur. To the extent that this conflict is known to both Margot and Dan, it may make this part of the planning process feel adversarial and tense. And of course, if these sources of adversariality are all taken to extremes, it becomes difficult to explain why interlocutors would engage with each other at all.[11]

Relevant conflicts in agents' unshared intentions can, in this way, create fric-

[11]It is perhaps worth noting here that in some of the most extreme examples of adversarial communication, such as cross-examination in court, we have to rely on incentives to cooperation, such as contempt-of-court and perjury laws, that function as artificial substitutes for the cooperative principle and Gricean maxims.

tion in the planning process. In this case, we're not talking about a conversation plan, but a regular old shared plan to host a dinner party. (Of course, it *is* a shared plan that the planners are constructing by means of conversation. But it isn't a plan about that very conversation.) But the same kind of friction can arise in conversation plans themselves, and this is part of our explanation of why adversarial conversations tend to restrict our communicative options. Consider again the following examples, from §1:

(9) A: Were you at the Tipsy Elf at noon yesterday?
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 a cooperative conversation, but not in at least some adversarial conversations, B's utterance in (9) would normally be taken as evidence that they were implicating that they hadn't been at the Tipsy Elf at any time close to noon. After all, we said, if they wanted to be cooperative and had been there any time close to noon, they probably would have volunteered this information, as in (10). But since they didn't do this, they must have intended A to infer that that they weren't there.

We can now make this point more explicit, and give a bit more of an explanation. The kind of cooperativity that is relevant to explaining this case consists in the coherence of A's and B's intentions. A is asking about B's whereabouts at noon yesterday for some reason—i.e., in order to further some antecedent plan. In many such cases, information about whether B was at the Tipsy Elf at half past noon would be just as relevant to furthering this plan as information about whether B was at the Tipsy Elf at noon. For example, A might be trying to express interest in B's day, or trying to find out about whether B might have run into a mutual friend who went to the Tipsy Elf for lunch, or—if A is interrogating B—trying to establish that B committed a crime at the Tipsy Elf yesterday. Suppose that B knows about, or can infer, A's reasons for asking the question. In this case, B might answer in a way that is responsive to A's reason for asking the question—i.e., to A's private superplan of the QUD—in addition to the question itself. But this responsiveness can take two forms. If B's own plans fit well with A's, then they might volunteer extra information in an effort to further these plans, above and beyond answering the question. This is what is happening in (10), and the absence of such behavior in (9), combined with the assumption that

B is cooperative, gives A good reason to infer that they can't think of any information that would help A to further their plans. By contrast, in the setting of an interrogation, in which B is asking the question as part of a plan to establish A's guilt, and B wishes to maintain their innocence, it would be very surprising for B to offer more information beyond a simple "no," and this expectation explains why A won't infer that B meant anything extra by their utterance.

8 Conclusion: Genre and Pragmatic Competence

Pragmatic competence is our label for the human capacity to perform and interpret speech acts in the normal way. Construed narrowly, this is the competence to perform actions with a suite of audience-directed intentions and to recognize others as doing the same. More broadly, however, this can involve a whole range of complex plans, beliefs, and desires, insofar as they feed into processes of speech production and interpretation.

We have argued that there are fundamentally distinct genres of conversation—constituted by different patterns in shared conversation plans—and we think, furthermore, that the acquisition of pragmatic competence involves becoming sensitive to these variations. Competent speakers acquire a cognitive facility with these basic genres, making the difference between, for instance, exploratory and committal interactions easily recognizable. This facility provides an effective tool for resolving problems in interpretation, making the production and interpretation of speech acts vastly easier than it otherwise would be. It also adds to the kinds of messages we can get across, especially indirect speech. Finally, this is also important to syntactic and semantic competence, since natural language includes mechanisms for manipulating and probing planning structures. One important piece of evidence for this, discussed in §4, is data on how prosodic focus is exploited to help interlocutors coordinate on the QUD. If we are right, however, and the QUD normally represents an important part of a larger conversation plan, then focus is really concerned with speakers' coordination on such plans and, ultimately, conversational genres.

Now, compare a genre-bound conversation with a highly plastic conversation—one in which there are few jointly accepted goals governing the conversation. For example, imagine that you have just walked up to an acquaintance at a party and although it is obvious that the two of you will have a conversation, it is not yet clear what you will discuss, whether your discussion will be serious or playful, factual or practical, and so on. In some ways, this type of conversation is freeing:

you are awash in possibilities. But in other ways it can be stressful. You might find yourself groping around for something to say, and you might find that more of what you say will be misinterpreted than in a more balanced conversation.

This scenario points to a coordination problem which genres are designed to solve. Sometimes we solve coordination problems with conventions or social institutions, but human conversations should be fertile ground to look for more basic cognitive mechanisms and recurring patterns of interaction. And this is exactly what we get if we think about conversations as governed by planning structures.

References

Beddor, R. and Egan, A. (2018). Might to better: Flexible relativism and the qud. *Semantics and Pragmatics*, 11(7):1–43.

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

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

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

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.

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

Cohen, P. R. and Levesque, H. J. (1985). Speech acts and rationality. In *23rd Annual Meeting of the Association for Computational Linguistics*, pages 49–60. Association for Computer Linguistics, Chicago.

Cohen, P. R. and Levesque, H. J. (1990). Intention is choice with commitment. *Artificial Intelligence*, 42:213–261.

Cohen, P. R. and Perrault, C. R. (1979). Elements of a plan-based theory of speech acts. *Cognitive Science*, 3:177–212.

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

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

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.

Ginzburg, J. (1995a). Resolving questions, i. *Linguistics and Philosophy*, 18(5):459–527.

Ginzburg, J. (1995b). Resolving questions, ii. *Linguistics and Philosophy*, 18(6):567–609.

Ginzburg, J. (2012). *The Interactive Stance: Meaning for Conversation*. Oxford University Press, Oxford.

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.

Groenendijk, J. and Stokhof, M. (1984). *Studies on the Semantics of Questions and the Pragmatics of Answers*. PhD Dissertation, Institute for Logic, Language, and Computation (ILLC), University of Amsterdam.

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

Hamblin, C. L. (1973). Questions in montague english. *Foundations of Language: International Journal*, 10(1):41–53.

Hoek, D. (2018). Conversational exculpature. *The Philosophical Review*, 127(2):151–196.

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

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

Lepore, E. and Stone, M. (2015). *Imagination and Convention*. Oxford University Press, Oxford.

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.

Perrault, C. R. (1990). An application of default logic to speech act theory. In Cohen, P. R., Morgan, J., and Pollack, M. E., editors, *Intentions in Communication*, chapter 9, pages 161–185. MIT Press, Cambridge, MA.

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.

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.

Schaffer, J. (2007). Knowing the answer. *Philosophy and Phenomenological Research*, 75(2):383–403.

Schoubye, A. and Stokke, A. (2016). What is said? *Noûs*, Online First.

Simons, M. (2001). Disjunction and alternativeness. *Linguistics and Philosophy*, 24:597–619.

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.

Simons, M., Tonhauser, J., Beaver, D., and Roberts, C. (2010). What projects and why. *Semantics and Linguistic Theory*, 20:309–327.

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. (1999). *Context and Content*. Oxford University Press.

Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*, 25:701–721.

Stojnić, U. (2018). Discourse, context and coherence: The grammar of prominence. In Preyer, G. and Peter, G., editors, *Beyond Semantics and Pragmatics*, pages 97–124. Oxford University Press, Oxford.

Thomason, R. H. (1990). Accommodation, meaning, and implicature: Interdisciplinary foundations for pragmatics. In Cohen, P. R., Morgan, J., and Pollack, M. E., editors, *Intentions in Communication*, chapter 16, pages 325–364. MIT Press, Cambridge, Mass.

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

von Stechow, A. (1991). Focusing and backgrounding operators. In Abraham, W., editor, *Discourse Particles: Descriptive and Theoretical Investigations*, volume 12 of *Pragmatics & Beyond: New Series*, pages 37–84. John Benjamins, Amsterdam.

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

Yalcin, S. (2018). Belief as question-sensitive. *Philosophy and Phenomenological Research*, 97(1):23–47.