Participants seek shared outlooks in non-canonical disagreements: Evidence from a corpus of dyadic conversation in English

John Duff and **Lalitha Balachandran**
Department of Linguistics
UC Santa Cruz
jduff, lalithab@ucsc.edu

Abstract

We present data and a preliminary analysis of a novel kind of disagreement observed in a corpus of English dyadic conversations. In conversations about a variety of topics, speakers volunteer attitude and speech reports rather than direct answers to Questions Under Discussion. The conversations are challenging to capture with formal pragmatic models, not just because of the mismatches between what is said and the QUD, but also because they lead to apparent disagreements despite a lack of contradiction. We analyze these cases as participants aiming to align non-committal stances, inspired by recent approaches that treat subjective conversation as the coordination of outlooks. Overall, the discussion advocates for flexible models of conversation that allow extra-linguistic goals and pressures to interface with the lower-level dynamics of discourse moves.

1 Introduction

The standard approach to modeling conversations between pairs of participants in formal semantics and pragmatics has involved the notion of *joint commitment*. Conversational participants make and invite assertions with the goal of growing the common ground, i.e. the list of discourse commitments held by both participants (Stalnaker, 1978; Farkas and Bruce, 2010; Farkas and Roelofsen, 2017). These models of conversation primarily deal with the dynamics of information exchange in discourse.

Sometimes participants' discourse commitments stand in the way of information exchange. *Canonical disagreements* like (1) lead to so-called "conversational crisis", seemingly requiring the participants to either retract some commitment or agree to disagree. Information exchange models provide an explanation for why this is: A and B's assertions cannot both be true, as they are directly contradictory, and so they stall the addition of information to the common ground.

(1) A: This tree is taller than it was yesterday.
B: No, it's not!

This paper examines conversations like (2)—intuitively, also a "disagreement", though it lacks direct contradiction. We might imagine that A and B will continue by trying to change each others' minds. But a simple information exchange model doesn't explain these intuitions: here, A and B both make true assertions, and it is unclear why this sequence of exchanges should stall the conversation if participants were simply informing one another of contrasting attitudes with no further consequence.

(2) A: I think this tree is taller.
B: Well, I don't!

Many models assume that discourse is organized according to, usually implicit, *Questions Under Discussion* (QUDs) (Roberts, 1996/2012; Büring, 2003). The QUD structure of a discourse captures patterns of coherence and relevance between participants' conversational moves. But intuitively, conversations do not always follow this idealized template. At times, higher-order conversational goals drive the way in which a particular sequence of discourse moves unfolds, leading to strategies of inquiry that are not directly tied to the QUD in form. This is what we will say about examples like (2), which we will call *non-canonical disagreement*.

In a case study of three dyadic conversations, we show that non-canonical disagreements arise and are resolved with the same basic signature, regardless of the subjectivity of the QUD guiding the discourse. First, participants establish explicitly autocentric viewpoints as in (2). If these viewpoints contrast, they embark on a longer process of justifying their decision process, and attempt to reach a joint outlook on the issue. This process resembles canonical disagreement despite the fact that participants nowhere establish an actual contradiction.

We ultimately propose that non-canonical disagreements (i) arise when participants embark on

Proceedings of the 26th Workshop on the Semantics and Pragmatics of Dialogue, August, 22-24, 2022, Dublin.

a somewhat divergent strategy of inquiry concerning questions of shared viewpoint and (ii) may be settled only if they establish parallel *stances* with respect to a proposition p_n via a process of inference. We will argue that both features might be fruitfully derived using outlook semantics (Coppock, 2018) within an augmented version of the Table model of conversation (Farkas and Roelofsen, 2017).

2 Methodology

Undergraduate students at UC Santa Cruz were recruited in 2020 and 2021. They participated remotely, and were compensated with course credit. First, participants privately reviewed four sets of art or media, each set presented with a free response question. Participants then joined each other in a video call, and were instructed to discuss the art and media they saw and the questions they answered for twenty minutes. They were not explicitly instructed to reach an agreement or record revised answers after discussion. Each participant's audio and video was recorded during this conversation, and an automatic aligned transcript was prepared using the text-to-speech service Descript.

Each set appeared with either an objective or subjective question.[1] For instance, one set included three photographs of Bruce Springsteen. Some dyads were asked the (objective) question "Who is depicted in these images?"; others were asked "Which is the worst picture of the musician?", subjective due to the multidimensional adjective *worst*.

We present here data from hand-corrected transcripts of three dyads, all native speakers of English.[2] Participants 11A (23 F) & 11B (19 F) were paired randomly, while participants 12A (19 F) & 12B (19 F) and 13A (19 F) & 13B (20 F) were friends who signed up to participate together.

We identified eleven disagreements in these transcripts, including both explicit contradictions and self-ascriptions of non-cotenable viewpoints. Six were disagreements about the experimental questions (four objective, two subjective). Another five were about questions that came up organically.

The authors annotated these eleven disagreements with the explicit QUDs introduced by the experiment, and implicit QUDs reconstructed such that participants' moves could be construed as intuitively relevant to the current QUD (see §4). The resulting QUD structures provided a framework to understand the organization of each conversation.

3 Data

This section establishes some basic descriptive generalizations from the three conversations, including many cases of what we call non-canonical disagreement. We show that non-canonical disagreements have a consistent profile: participants self-ascribe differing attitudes or judgments, and, finding themselves in a dispute, take turns justifying their positions with the goal of reaching a joint viewpoint.

3.1 How non-canonical disagreements begin

In discussion of the QUDs, dyads typically began by self-ascribing attitudes towards potential answers to the explicit QUDs. We focus here on cases where those attitudes differ, and discuss how these conversations continue in §3.2.

In this initial stage, participants often self-ascribed attitudes in varying tenses: for instance, in (3), 11A reports a past attitude in order to establish a contrast with a present attitude of 11B. When using the past tense, speakers seemed to be speaking of their attitudes as they viewed the stimuli and answered the given questions.

(3) QUD: *Were these (three) videos produced in the same decade?*

11B: ...**number three had some like animation** at the same time, but **I think that's pretty, that's like more advanced**...

11A: Interesting. See now ... I also do not think they're the same decade, but **I thought that clip two was actually the newest**.

11B: Oh.

Participants also very frequently spoke only of their answers, in which case their attitude towards the propositions that formed their answer is expressed only indirectly. Nevertheless, their attitudes still seem to be the main point of such assertions: e.g. in (4), it is perfectly coherent for 13B to respond to 13A's description of an answer by describing a (contrasting) past-tense attitude.

[1] Throughout this paper, we use the word SUBJECTIVE to pick out all and only the kinds of content about which "faultless disagreement" seems to be licensed (Kölbel, 2004, see also §5.1). Subjective questions thus included predicates of personal taste like *tasty* and *beautiful* (Lasersohn, 2005), but also multidimensional adjectives like *good* (Sassoon, 2013).

[2] Full transcripts of these conversations are available at https://osf.io/jwye8. While the other 47 transcripts in the corpus remain to be corrected, we anticipate making the full corpus publicly available for future research.

(4) QUD: *How many people are singing in each of these recordings?*

13A: **I just put**...one person, in the first one, **two in the second one**, maybe three or more in the last one.

13B: Okay, **I thought the second one only had one person**.

In the above cases, the participants' opening assertions are unprompted, but in other cases they were guided by explicit questions, as dyad 13 demonstrates in (5). In (6) we can see a similar case: 13A asks 13B for their attitude. 13B's response, a description of what they wrote during the task, shows another example of verbs like *put* serving as indirect attitude ascriptions.

(5) QUD: *Which artwork is the most beautiful?*

13B: Which one is most beautiful? Which one did you put?

13A: **The first one**.

13B: Oh, really?

13A: What'd you put, the third one?

13B: No, **I put the second one**...

(6) QUD: *Which is the best picture of Alcatraz?*

13A: Which picture did you like better?

13B: Oh, **I put number two**. What about you?

13A: **The first one**.

Note that these patterns of self-ascription were the most common opening for every experimental question we examined, often producing cases where participants found their attitudes aligned (7).

(7) QUD: *When were these films made?*

12B: I said like nineteen forty.

12A: Okay. Yeah. I was ... in that ballpark.

We also observe no clear differences in behavior across objective QUDs like those in (3-4) vs. subjective QUDs like those in (5-6), featuring the PPT *beautiful* and the multidimensional superlative *best*.

3.2 How non-canonical disagreements settle

In the examples above, participants use first-person indexicals to contribute their respective answers to the QUD without directly disagreeing. They never dispute the accuracy of one another's responses. Nevertheless, when they discovered they held differing viewpoints, they entered into lengthy discussions concerning the validity of each participants' viewpoints, beginning by justifying their own positions and then arguing against their interlocutor's. In this process, agreement was often sought, and freely given on matters of simpler evidence or taste.

For instance, early in the process of resolving the disagreement in (5), 13A directly critiques the piece of art which 13B listed as *most beautiful* (8). Note that 13B agrees with 13A's observations without conceding the validity of their own judgment, which they continue to defend after this excerpt.

(8) (QUD: *Why didn't 13A pick the second?*)

13A: I just thought the second was kind of ... it's like mostly the same color.

13B: Yeah, you're right. ... I had to like really look at it to see what it was.

13A: Yeah, like there's no blacks, or like dark, dark colors.

13B: Hm.

13A: Yeah, there's no dark colors.

13B: Yeah, that's true.

Similarly, after (3), 11A goes on to explain the basis of their position regarding the age of dyad 11's video clips. Again, though 11B here accepts 11A's observations (*mhm* and *yeah*), they continue later in the conversation to contest whether 11A's answer is appropriate.

(9) (QUD: *Why did 11A think clip two was the newest?*)

11A: But ... clips one and three had more of that. Like, you know, that graininess, um—

11B: Yeah.

11A: ...which kind of strangely makes me believe that that was produced much, uh—I said later, but later in the time period, so.

11B: Mhm ... yeah.

Participants do not always seek to justify their own position and levy critiques at others; they can explicitly cooperate in their interlocutor's justification. For instance, in the process of resolving the disagreement in (6), 13B sympathizes with 13A's attitudes towards 13A's favored photo, and later invites them to further explain their preference.

Likewise, the participants also seek out points of agreement on questions that we might think are

properly outside of the scope of the disagreement at hand. In dyad 13's same dispute regarding photographs, shortly after each participant has laid out their opinions of each others' favored photos, 13B seeks out a joint opinion about the third photo, which no one chose (10).[3]

(10) 13B: Can we agree that the last one was terrible?

13A: Yeah, I didn't like that one. {LG} There's like a, um, a pole in the middle of it.

As these processes of justification and alignment continue, it is apparent that participants are aiming to avoid an outcome where they maintain different viewpoints. For instance, dyad 13 doesn't move on from the *best picture* and *beautiful art* conversations until they seem to have reached a consensus. The participants reflect explicitly on the pressure they feel to do this (11).

(11) 13A: Okay, **I guess technically the best picture is the second one**.

13B: Okay. Thank you for caving.

Other times, consensus seems to be reached very easily. In (12), after (4), in the face of 13B's argument, 13A readily changes their mind.

(12) (QUD: *How many people are singing in the second recording?*)

13B: ...for the second one, I'm pretty sure they're harmonizing too.

13A: **Okay.**

13B: So I think there's more than one...

13A: I was thinking of instruments. {LG}

But consensus isn't always reached, and in cases of apparently intractable disagreement, participants sometimes moved on without reaching a joint viewpoint. In such cases participants closed the conversation by re-affirming their different attitudes, as dyad 13 do in (13).

(13) (QUD: *Which piece of music is the best?*)

13B: Maybe **I would say the third one** then.

13A: You'll say the third one?

13B: Yeah.

13A: Okay. **I still like the first one**.

13B: Okay, cool.

[3]The annotation {LG} indicates speaker laughter.

This happened even for non-canonical disagreements for objective QUDs, as in (14). After a protracted dispute about the identity of a musician in a series of photographs, 12A uses a reverse image search (visible only to 12A) to obtain evidence, and reports back that they are convinced the man depicted is Bruce Springsteen. 12B nevertheless remains doubtful, and while 12A acknowledges that it is possible the photos depict someone else (*maybe you're right*), the dyad concluded their conversation with very different apparent belief states.

(14) (QUD: *Is this Bruce Springsteen?*)

12B: I don't think it's Bruce Springsteen.

12A: It so is. It came up. (*in the search*)

12B: I, I don't trust it then.

12A: Okay, two of them (= *pictures*) came back saying Bruce Springsteen.

12B: ...Something in my bones is saying ... it's not that. And I am not a Bruce Springsteen expert, but just–

12A: **I don't know, maybe you're right.** Maybe like someone dressed up and tried to impersonate him or something.

12B: ...**I almost don't want to know who he is**.

3.3 Interim summary

Across the three conversations, we see examples of disagreements with the same general properties. They begin with the establishment of autocentric viewpoints. When viewpoints contrast, a longer discussion ensues where participants review their evidence and decision-making processes. The goal of this process seems to be to negotiate which viewpoint they should collectively adopt, each trying to either collect enough evidence to change their own mind or present enough of an argument to change their interlocutor's. When this goal is achieved successfully, i.e. they reach congruent viewpoints, the QUD is notionally resolved. When they fail to reach a consensus, they simply agree to disagree, as with any canonical disagreement where resolution is not successful. These properties held for objective and subjective questions alike.[4]

In the remainder of the paper, we aim to understand why participants establish autocentric view-

[4]We further note that the same patterns show up on a cursory examination of other conversations in the corpus, across participants regardless of gender and age.

points, and why they aim to agree on a shared viewpoint. This behavior is unexplained in basic models of discourse as information exchange: autocentric viewpoints are not appropriately relevant to the apparent QUD, nor are they sufficient commitments to resolve it, nor are contrasting viewpoints clearly problematic in any way. We will propose that these discrepancies arise because participants are making two non-canonical choices in the structure of their conversation, potentially due to insufficient evidence to fully settle the QUD: (i) they follow a strategy of inquiry which is indirectly related to the QUD, and (ii) they transmute the QUD itself into one which is resolved by a joint outlook rather than joint objective commitments.

4 The relevance of establishing viewpoints

4.1 QUD preliminaries

Strict versions of QUD theory maintain that an assertion must be *relevant* to the QUD that dominates it (Roberts, 1996/2012). Given the objective QUD in (7), *When were these films made?*, participants may be relevant by asserting one of the full answers {*They were made in the thirties, They were made in the forties, ...*}. Alternatively, they may project a *strategy of inquiry*, which requires pursuing answers to subquestions (e.g. *When was film {1,2...} made?*) that are entailed by the higher-level QUD .

One advantage of standard QUD theory is that it captures the information structural relationship between assertions in a discourse and their corresponding QUD structures, which is mediated via *focus*. For each proposition in a set of answers to a question, the focus is associated with alternative expressions (*the thirties, the forties, ...*), whereas the backgrounded content remains constant. The assertion-QUD correspondence, then, is often assumed to be fairly direct.

4.2 Establishing viewpoints

The data presented in §3 constitute a puzzle for models of conversation that adopt standard QUD theory: why is it that a pair of autocentric assertions (e.g. *I said nineteen forty* and *Yeah, I was in that ballpark* in the context of the QUD $p?$: *When were these films made?*) contributes information that is treated by the participants as relevant to the overarching QUD? Note that $p?$ does not entail a subquestion about A or B's attitudes. Nevertheless, exchanges such as the one in (7) seem to be coherent, and even successful in addressing the QUD.

There are similar cases where assertions have been argued to indirectly correspond to their QUDs. For instance, Simons (2007) examines examples like (15), where a proposition embedded under a reportative or attitude predicate exhibits a so-called *embedded main point use*.

(15) A: What's the weather like?
B: Jane said that it's raining.

Here, B's embedded content directly answers the QUD, whereas the matrix content serves to provide information concerning the availability or quality of evidence for the embedded proposition, but only provides an indirect answer to the QUD. If B's response were a direct answer, we would expect the QUD, roughly, to be: *What did Jane say (about the weather)?* Note that this question is not entailed by A's question in (15). Nevertheless, the intuition is that B's response is coherent.

The cases examined in §3 take the same general shape as in (15). That is, participants' assertions are not directly relevant to the explicit QUD.

(16) $p?$: *When were these films made?*
A: I said nineteen forty.
$\rightsquigarrow q?$: (*When did A say they were made?*)

(17) B: I said nineteen fifty.
$\rightsquigarrow r?$: (*When did B say they were made?*)

Note that subQUDs $q?$ and $r?$, too, cannot be part of a strategy of inquiry in the sense of Roberts (1996/2012), as they are not entailed by $p?$. To treat them as nevertheless coherent, we may adopt more relaxed constraints on relevance and entailment, following Riester (2019) and others, though we will have to say more about how exactly they satisfy the interlocutors' purposes.

Concretely, we propose that when participants are faced with a QUD that they have insufficient evidence to address directly,[5] they may choose to adopt a mediating strategy of inquiry which involves the projection of individual autocentric, attitudinal subQUDs, e.g. *When did {A,B} say/think they were made?* in (16). Their choice may be driven by competing conversational pressures: here, complying with the Maxim of Quality may override the pressure to maintain relevance. Given the initiation of this strategy by a speaker, the responding participant is likely to adhere to the same

[5] Notably, the one canonical disagreement in these conversations, where participants make full assertions rather than self-ascribing attitudes, is one about a topic which participants clearly have solid prior knowledge.

strategy (i.e. respond by merely establishing a viewpoint), unless they are more confident in the quality of their own evidence.

The proposed strategy of inquiry explains the form and coherence of the participants' assertions within standard QUD theory, but the question remains how this strategy of inquiry is itself relevant to the QUD dominating it. Our intuition is that a conversation about individual viewpoints serves to allow the participants to assemble evidence and arguments towards an answer to the QUD. But we aim also to explain why this process is not typically complete until participants establish parallel attitudes. In §5, we propose that in addition to adopting this strategy of inquiry, participants also alter their overarching objective, shifting from the search for concrete information to the goal of aligning their hypotheses about the world.

4.3 Retractions

In a situation where one participant manages to convince the other to change positions with respect to their attitude on $p?$, we do not see explicit retraction of the compromising participant's original attitude. That is: explicit, linguistically identifiable acknowledgement of a change in commitment, e.g. *I was wrong*, is non-existent in our data set.

This is perhaps due to the fact that retraction is actually unnecessary. One advantage of asserting an attitude with respect to $p?$, rather than directly addressing the issue of $p?$, is that participants are able to concede their original viewpoints without retracting, as they have only committed to an attitude at a particular time. That is, if the proposition *A thinks the first picture is the best one at time* t is in the common ground and A adds the proposition *A guesses at time* t' *that the second picture is technically the best one*, this expresses a change in A's attitude, but A hasn't made contradictory commitments. Given uncertainty about the answer to a QUD, this is a useful strategy, as it allows participants to assert their attitudes without needing to resolve possible later self-contradictions.

5 Resolving non-canonical disagreements

In the previous section, we suggest that these conversations involve strategies of inquiry that project subQUDs about everyone's attitudes. This allows us to capture the local relevance of those viewpoint-establishing moves, but two puzzles remain. First, how can these attitudes satisfy the participants' goals for the conversation? And second, what is different about non-canonical disagreements such that they don't satisfy those goals?

In this section, we review previous approaches to settling QUDs, and propose that participants in these conversations are actually settling a QUD of a non-transparent form, akin to the QUDs in conversations about taste. The proposal can account for how these conversations get settled without assuming that participants have reached an agreement on the actual state of the world.

5.1 Question resolution in discourse

In modern commitment-based discourse-models, when a QUD is on the table,[6] participants cannot treat that QUD as settled until it has been *resolved*.

(18) **QUD Resolution:** A QUD $p?$ is resolved iff participants have collectively committed to one of its possible answers p_n.[7]

On this approach, one of the principal features of a canonical disagreement is that QUD resolution is blocked unless someone retracts one of their commitments. While a QUD remains unresolved, participants must continue working to establish an answer, or else give up the search a joint answer, perhaps engaging in meta-linguistic negotiation to remove or change the QUD (Ginzburg, 2012).

Here lies the problem with non-canonical disagreements: the viewpoint-establishing moves do not generate the commitments needed to resolve the QUD, even when participants establish the same viewpoint. Likewise, even when viewpoints are not aligned, they do not block QUD resolution.

Similar discrepancies between participants' utterances and their discourse effects are at-issue in work on subjective meaning. Consider (19).

(19) A: This chili is tasty.
  a. B: Yes, it is.
  b. B: No, it's not.

One influential approach to subjectivity since Lasersohn (2005) offers a relativist semantics in which predicates like *tasty* might be true or false of the same tasted object within the same world depending on the identity of a judge parameter in the context of assertion. A's assertion would only commit

[6] Classically, table models (Farkas and Bruce, 2010; Farkas and Roelofsen, 2017) track only explicit QUDs. We will assume here that implicit QUDs can also enter the table; see Ginzburg (2012) for a similar proposal.

[7] For Farkas and Roelofsen (2017), once the set of worlds compatible with everyone's commitments entails p_n.

A to an autocentric judgment (see Stephenson, 2007). But on a relativist semantics, three puzzles remain for the pragmatics: (i) why does A's move seem to project a more general QUD about the tastiness of the chili?; (ii) in (19a), why does that QUD seem to be settled?; and (iii) in (19b), why does there seem to be pressure to continue on the same topic until a state like (19a) is reached?

A compelling response to these puzzles that has emerged in the ensuing pragmatic literature is that interlocutors in these conversations are collaborating not to narrow down the set of possible worlds they collectively might inhabit, but the set of what Coppock (2018) calls *outlooks*, refinements on possible worlds that include positions on subjective issues (see also Stephenson, 2007; Egan, 2010; Rudin and Beltrama, 2019). In an outlook-based semantics, we can speak of truth or falsity of a given proposition in the actual outlook of the author of an information state. In a commitment-based conversational model built on top of this, participant's individual discourse commitment sets describe the outlooks they represent themselves as having, while the common ground encompasses joint outlooks. In particular, this accounts for the above intuitions without giving up on the classic intuition of "faultlessness" in disagreements like (19b) (Kölbel, 2004; MacFarlane, 2014): both participants are asserting felicitously based on their outlook, even though by doing so they ultimately block resolution of the QUD.

We will assume this formalism for subjective meaning. As described in §3.2, in the disagreements we are discussing, the profile of participants' reactions to either type of disagreement was largely the same. For these reasons, we will mostly abstract away from the differences between objective and subjective expressions in what ensues.

5.2 Pondering: When stances are enough

Basic models of conversation as information exchange cannot capture why stances about answers to a QUD appear to settle it or prevent settling it. We propose instead that participants in these conversations are not aiming to resolve the questions that were posed to them *per se*, but instead aiming to reach a type of joint outlook on those questions, employing the formalism of Coppock (2018).

In the spirit of conversational models like Ginzburg (2012) that outline the ways participants might negotiate changes to the current parameters of their conversation, we suggest that participants have at their disposal a conventional parameter change procedure we'll call PONDER.

(20) PONDER: When participants think they cannot adequately answer a QUD $p?$ with answers $\{p_1, p_2 ... p_n\}$, they may replace q with an alternative QUD $p'?$ with answers $\{p'_1, p'_2 ... p'_n\}$ such that p'_n is true for a given outlook iff that outlook includes a positive stance towards p_n.

(21) **Positive stance**: An individual has a positive stance towards p_n in the context of a QUD $p?$ iff among the answers of $p?$ they are most willing to entertain that p_n is the case.

We'll call the QUDs PONDER generates *stance-QUDs*. The answers to these stanceQUDs are discretionary propositions as defined by Coppock (2018), which concern the views of the individual(s) committed to them, and can be true or false in a given world, depending on the outlook in question. We might paraphrase a stanceQUD as "Which answer p_n to $p?$ are we most willing to entertain?", though crucially they are dependent not on deictic *we* but a Lasersohnian judge. Lacking evidence of the explicit form of stanceQUDs (or whether a form exists), we might adopt a somewhat liberal position on the nature of QUDs: while there is often a natural correspondence between QUDs on the table and their syntactic form, perhaps this need not always be the case.[8] Minimally, stanceQUDs are abstract goals with resolution conditions that we can represent formally.

To settle the stanceQUD, participants project a strategy of inquiry that is not directly related in form, as discussed in §4. This strategy of inquiry is what invites participants to assert e.g. first-person attitudes, and it is only through inference that these assertions are taken to establish a stance, rather than merely commit to a first-person attitude. It is thus a special pragmatic effect of such assertions in the contexts we are discussing that they also enter an answer to a stanceQUD into the speaker's discourse commitments.[9] Because this approach is

[8] This flexibility would make structural relationships (e.g. Büring, 2003) between stanceQUDs and subQUDs impossible, but note that we have already relaxed that assumption in §4.

[9] First-person assertions live a similar double life with canonical subjective QUDs. In (i), the QUD seems settled, but with a Coppock-style representation, the objective assertions could not directly resolve it. As in our cases, if we take the assertions to imply certain outlooks, we can see why the QUD has been resolved.

somewhat novel and formally complex, we demonstrate in detail how it would model non-canonical agreement and disagreement in Appendix A.

If both speakers indirectly establish p'_n—a positive stance towards answer p_n of the original QUD—they will have formally resolved the stance-QUD, capturing why these conversations seem to be settled. In contrast, if anyone infers a stance for a speaker that differs from their interlocutors', a conversational crisis will ensue. We could assume that participants refrain from inferring stances in this case, but if so, the stanceQUD remains unresolved. To resolve it, participants must continue discussion of the matter until both are willing to establish the same stance. In this way, we capture the pressure for stance alignment observed in §3.2 as a species of the same pressure observed in any conversation, to answer all QUDs on the table.

5.3 Resolving the QUD via inference?

We have proposed that participants are able to implicitly alter the structure of a QUD to introduce a goal with properties more similar to a subjective question, that is resolved through a joint stance. One can imagine another analysis: instead of assuming an implicitly altered QUD, why not assume implicitly strengthened commitments? We'll give one argument against the latter approach.

For viewpoint-establishing moves to resolve the apparent QUD $p?$, it would have to be the case that participants infer that p_n is part of a speaker's discourse commitments when that speaker expresses a positive view of p_n. This is not *prima facie* unreasonable: consider the premise in (22).

(22) **Commitment to attitudes:** A participant in a conversation where p_n is relevant with a positive view of p_n should be committed to p_n for the purposes of the conversation.

With this premise, parallel viewpoint-establishing answers to personalized subQUDs would make participant commitments about the main QUD readily inferable. Once all participants have established a positive view of p_n, p_n will be assumed to be part of all of their discourse commitments, and thus the QUD can be resolved.

But (22) seems to crucially mischaracterize what we usually infer when participants self-ascribe a viewpoint. The argument against it follows the objection Simons (2007) raises against treating cases like (15) as "assertive" (Hooper, 1975). If in (15) B is understood to commit that it's raining, we could understand how their assertion answers A's question. But this analysis misses another classic Gricean implicature of B's utterance, that by avoiding a more direct locution, B gives the impression that they are unwilling to assert that it's raining.[10] The same critique is relevant here. Speakers who merely establish a positive view of p_n are specifically and effortfully avoiding full commitment to p_n. It runs counter to that avoidance to assume they are implicitly committing to p_n.

In contrast, the stanceQUD account manages to capture the ways in which establishing a stance settles a QUD, without dangerously assuming that all participants are representing themselves as committed in full to a particular answer. It is ultimately an empirical question whether participants take stances as evidence for implicit commitments, but until such evidence can be established, we take our proposal to be preferable.

6 Extensions and upshots

6.1 Predicting (in)felicitous responses

We briefly note one piece of evidence to support the validity of the subQUD structure that governs autocentric strategies of inquiry.

(23) $p?$: *(What about It?)*

12B: I don't know if it's weird but I just got like slight *Devil All The Time* vibes from *It*.
$\rightsquigarrow q?$:(*Did B get Devil All The Time vibes?*)

12A: Yes!
$\rightsquigarrow r?$: (*Did A get Devil All the Time vibes*?)

In (23), A responds with the positive polar response particle (PRP) *Yes*. This leads to the "sloppy" interpretation that A also got these vibes from the movie *It*, as opposed to the strict interpretation, which would simply affirm B's assertion. Moreover, this seems to be a general property of PRP responses to attitudes: responding *Yes* to *I hope it rains tomorrow* can only mean that the responding participant also hopes that it will rain. In contrast, a positive PRP response to a non-attitude report such as *I had a bad dream* is infelicitous. The fact that

(i) QUD: *Is this chili tasty?*
A: I like this chili.
B: I do too.

[10] See also Simons (2019) for a more recent argument against a relevance implicature analysis of (15).

PRPs can reference subQUDs $q?/r?$ provides evidence for this subQUD structure in conversations about attitudes, and suggests a way for future work to provide empirical tests for our claims here.

6.2 Subjectivity

Our proposal suggests there are two routes to explain cases of so-called "faultless disagreements." In addition to assertions which are properly judge-dependent, we have argued that stance self-ascriptions can be used to faultlessly disagree with regard to an implicit subjective QUD while being strictly objective in form. Faultless disagreement has been advanced as a diagnostic for the presence of relative truth, not just for predicates of personal taste and their ilk, but also epistemic modality (Stephenson, 2007; MacFarlane, 2011; see Weatherson and Egan, 2011) and statements about the future (MacFarlane, 2003; Giannakidou and Mari, 2018), even though the latter cases fail other diagnostics like *find*-embedding (Coppock, 2018). It's possible that impressions of disagreement for some of these cases come about not through the presence of bona fide judge-dependent meaning, but because in context they are being used to address implicit questions that require joint outlooks. We hope that future work, especially examining naturalistic conversations, might follow up on this possibility.

We also note that the similarities between participants' treatment of objective and subjective QUDs are good evidence for theories of subjective meaning that predict misaligned outlooks to be just as dire as incompatible commitments. We plan to continue looking for differences in behavior on a larger scale as we prepare the complete corpus.

6.3 Outlook congruence

We suggest, tentatively, that the desire for participants to reach a joint outlook may be driven by a general pressure to achieve social cohesion with one's interlocutor (Edwards and Middleton, 1986; Egan, 2010; Coppock, 2018). While not a requirement, this pressure would explain the preference to attempt alignment before leaving the QUD unresolved. The source of this non-linguistic pressure and its empirical validity remain somewhat underexplored, but this idea is consistent with other work on socially-induced QUDs in similar autocentric conversational contexts. For example, Balachandran (2021) argues that a social principle called the *Norm of Reciprocity*, which underlies a pressure for participants to reciprocate in situations involving avowals and conflicts, induces a QUD structure that has the ability to mediate instances of mismatching indexical reference in verb phrase ellipsis (see Chung (2000) and Charnavel (2019) for more detail).

(24) QUD: (*Do A and B love each other?*)

A: I_A love you_B.
$\rightsquigarrow$ subQUD: (*Does A love B?*)

B: Well, I_B don't <love you_A>!
$\rightsquigarrow$ subQUD: (*Does B love A?*)

In (24), A's assertion is taken to project an implicit QUD structure and compel B to respond to the subQUD *Does B love A?* The fact that B's response (24) appears to trigger disagreement is derived pragmatically: violation of the Norm of Reciprocity is taken to lead to interpersonal conflict, but does not block QUD resolution. In contrast, aligned stances are required to settle the QUD under the current analysis. Though both cases aim to derive a pressure for "alignment", here we enshrine this as a proper condition on QUD resolution. This is perhaps desirable, as the nature of these misalignments intuitively seem distinct in some sense, despite their similarities on the surface. Future work should aim to more thoroughly consider a pragmatic analysis of aligning stances.

6.4 Summary

In this article we have provided a description of a conversational phenomenon that proves challenging to treat using the basic toolbox of commitment-based discourse modeling. We suggested adding to that toolbox in two ways to account for these conversations: (i) allowing for implicitly projected strategies of inquiry that are not directly relevant to the current QUD, and (ii) formalizing how participants might pursue a shared hypothesis rather than a complete answer to a QUD. With these components in place, non-canonical disagreements look much like subjective disagreements, raising questions for future work on subjectivity and the role(s) of generalized social alignment in linguistic theories of discourse.

7 Acknowledgements

Jennifer Wheelock contributed to the design of the experiment in which these conversations took place, and assisted in data collection together with Dasha Komissarchik and David Tuffs. Robert Liu, Stephanie Palomares, Vindhya Shigehalli, Jacqueline Wen, and Naomi Wong assisted with transcription and annotation. The authors would like to thank Pranav Anand and Maziar Toosarvandani for helpful comments.

References

Lalitha Balachandran. 2021. Reciprocal questions and the pragmatics of argument reversing verb phrase ellipsis. In *Proceedings of Sinn und Bedeutung 26*.

Daniel Büring. 2003. On D-trees, beans, and B-accents. *Linguistics and Philosophy*, 26(5):511–545.

Isabelle Charnavel. 2019. Supersloppy readings: Indexicals as bound descriptions. *Journal of Semantics*, 36(3):453–530.

Sandra Chung. 2000. Close encounters with pronouns in VP ellipsis. In Sandra Chung, Jim McCloskey, and Nathan Sanders, editors, *Jorge Hankamer's WebFest*.

Elizabeth Coppock. 2018. Outlook-based semantics. *Linguistics and Philosophy*, 41:125–164.

Derek Edwards and David Middleton. 1986. Joint remembering: Constructing an account of shared experience through conversational discourse. *Discourse Processes*, 9(4):423–459.

Andy Egan. 2010. Disputing about taste. In Richard Feldman and Ted A. Warfield, editors, *Disagreement*, pages 247–286. Oxford.

Donka F. Farkas and Kim B. Bruce. 2010. On reacting to assertions and polar questions. *Journal of Semantics*, 27:81–118.

Donka F. Farkas and Floris Roelofsen. 2017. Division of labor in the interpretation of declaratives and interrogatives. *Journal of Semantics*, 34.

Anastasia Giannakidou and Alda Mari. 2018. A unified analysis of the future as epistemic modality: The view from Greek and Italian. *Natural Language & Linguistic Theory*, 36:85–129.

Jonathan Ginzburg. 2012. *The Interactive Stance: Meaning for Conversation*. Oxford.

Joan B. Hooper. 1975. On assertive predicates. In John P. Kimball, editor, *Syntax and Semantics*, volume 4, pages 91–124. Academic Press.

Max Kölbel. 2004. Faultless disagreement. *Proceedings of the Aristotelian Society*, 104:53–73.

Natasha Korotkova. 2016. Disagreement with evidentials: A call for subjectivity. In *Proceedings of SemDial 20 (JerSem)*, pages 65–75.

Peter Lasersohn. 2005. Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy*, 28:643–86.

John MacFarlane. 2003. Future contingents and relative truth. *The Philosophical Quarterly*, 53(212):321–336.

John MacFarlane. 2011. Epistemic modals are assessment-sensitive. In Andy Egan and Brian Weatherson, editors, *Epistemic modality*, pages 144–178. Oxford.

John MacFarlane. 2014. *Assessment Sensitivity: Relative Truth and its Applications*. Oxford.

Arndt Riester. 2019. Constructing QUD trees. In *Questions in Discourse*, volume 2, pages 164–193. Brill.

Craige Roberts. 1996/2012. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics & Pragmatics*, 5:6.

Deniz Rudin and Andrea Beltrama. 2019. Default agreement with subjective assertions. In *Proceedings of Semantics and Linguistic Theory (SALT) 29*, pages 82–102.

Galit W. Sassoon. 2013. A typology of multidimensional adjectives. *Journal of Semantics*, 30:335–380.

Mandy Simons. 2007. Observations on embedding verbs, evidentiality, and presupposition. *Lingua*, 117:1034–1056.

Mandy Simons. 2019. The status of main point complement clauses. In *Proceedings of SemDial 23 (LondonLogue)*, book section 15.

Robert C. Stalnaker. 1978. Assertion. In Peter Cole, editor, *Syntax and Semantics*, volume 9, pages 78–95. Academic Press.

Tamina Stephenson. 2007. Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30(4):487–525.

Brian Weatherson and Andy Egan. 2011. Introduction: Epistemic modals and epistemic modality. In *Epistemic modality*, pages 1–18. Oxford.

A Modeling stanceQUDs and their resolution

More or less in keeping with Farkas and Roelofsen (2017),[11] we take the state of a dyadic conversation to have at least the following components:

(25) TABLE: A stack of QUDs, represented as sets of answers.

(26) DISCOURSE COMMITMENTS ($DC_{A/B}$): The set of propositions each speaker is committed to for the purposes of the conversation.

(27) COMMON GROUND (CG): $DC_A \cap DC_B$, the set of all propositions the participants share joint commitments to.

We also assume the following:

(28) COMMITMENT SETS ($CS_{A/B}$): For a participant n, $\bigcap DC_n$, the set of all outlooks consistent with their commitments.

(29) CONTEXT SET (CS): $CS_A \cup CS_B$, the set of all outlooks consistent with all participants' commitments.

(30) **QUD Resolution** (formal):
A QUD $p?$ can be removed from the TABLE when $\exists p_n \in p?\,.\,CS \subset p_n$, that is, when participants' commitments entail an answer.

Now consider the toy conversation in (31).

(31) QUD: *Which video is the newest?*

A: I think the first video is the newest.

B: I also think the first video is the newest.

To model the conversation, we'll make reference to the following propositions and set of possible outlooks U.

(32) a. p_n is the set of outlooks where video n is the newest

b. q_n is the set of outlooks where A thinks video n is the newest at time t.

c. r_n is the set of outlooks where B thinks video n is the newest at time t'.

d. p'_n is the set of outlooks which include a positive stance towards p_n

(33) a. $U = \{w_1o_1, w_1o_2, w_2o_1 ... w_8o_2\}$

b. $p_1 = \{w_1o_1, w_1o_2, w_2o_1, w_2o_2, w_3o_1, w_3o_2, w_4o_1, w_4o_2\}$; $p_2 = U \setminus p_1$

c. $q_1 = \{w_1o_1, w_1o_2, w_2o_1, w_2o_2, w_5o_1, w_5o_2, w_6o_1, w_6o_2\}$; $q_2 = U \setminus q_1$

d. $r_1 = \{w_1o_1, w_1o_2, w_3o_1, w_3o_2, w_5o_1, w_5o_2, w_7o_1, w_7o_2\}$; $r_2 = U \setminus r_1$

e. $p'_1 = \{w_1o_1, w_2o_1, w_3o_1, w_4o_1, w_5o_1, w_6o_1, w_7o_1, w_8o_1\}$; $p'_2 = U \setminus p'_1$

We assume the conversation starts as follows:

(34) TABLE $= [\{p_1, p_2 ...\}]$
$CS_A, CS_B = \{w_1o_1 ... w_8o_2\}$
$CS = \{w_1o_1 ... w_8o_2\}$

The participants decide to apply PONDER:

(35) TABLE $= [\{p'_1, p'_2 ...\}]$
$CS_A, CS_B = \{w_1o_1 ... w_8o_2\}$
$CS = \{w_1o_1 ... w_8o_2\}$

They then project a strategy of inquiry which first involves the subQUD *Which video does A think is the newest?*. A's assertion adds the commitment q_1 to DC_A, and B automatically adds q_1 to DC_B, as A is an expert on A's attitudes (see Korotkova, 2016).

(36) TABLE $= [\{q_1, q_2 ...\}, \{p'_1, p'_2 ...\}]$
$DC_A, DC_B = \{q_1\}$
$CS_A, CS_B = \{w_1o_1 ... w_2o_2, w_5o_1 ... w_6o_2\}$
$CG = \{q_1\}$
$CS = \{w_1o_1 ... w_2o_2, w_5o_1 ... w_6o_2\}$

After (36), the subQUD is resolved and can be removed, because $CS \subset q_1$. The participants can also jointly infer from A's assertion that A has a positive stance towards p_1. This positive stance p'_1 is entered into A's discourse commitments, in turn restricting the outlooks the worlds that remain in their commitment set.

(37) TABLE $= [\{p'_1, p'_2 ...\}]$
$DC_A = \{q_1, p'_1\}$
$DC_B = \{q_1\}$
$CS_A = \{w_1o_1, w_2o_1, w_5o_1, w_6o_1\}$
$CS_B = \{w_1o_1 ... w_2o_2, w_5o_1 ... w_6o_2\}$
$CG = \{q_1\}$
$CS = \{w_1o_1 ... w_2o_2, w_5o_1 ... w_6o_2\}$

The next subQUD in the strategy of inquiry is *Which video does B think is the newest?*. B's assertion adds the commitment r_1 to DC_B, and A automatically follows suit.

[11] We crucially allow implicit QUDs on the table.

(38) $\text{TABLE} = [\{r_1, r_2...\}, \{p'_1, p'_2...\}]$
$DC_A = \{q_1, r_1, p'_1\}$
$DC_B = \{q_1, r_1\}$
$CS_A = \{w_1o_1, w_5o_1\}$
$CS_B = \{w_1o_1, w_1o_2, w_5o_1, w_5o_2\}$
$CG = \{q_1, r_1\}$
$CS = \{w_1o_1, w_1o_2, w_5o_1, w_5o_2\}$

After (38), the subQUD is resolved and can be removed, because $CS \subset r_1$. The participants can also jointly infer from B's assertion that B also has a positive stance towards p_1. This positive stance p'_1 is entered into B's discourse commitments, in turn restricting the outlooks the worlds that remain in their commitment set.

(39) $\text{TABLE} = [\{p'_1, p'_2...\}]$
$DC_A = \{q_1, r_1, p'_1\}$
$DC_B = \{q_1, r_1, p'_1\}$
$CS_A = \{w_1o_1, w_5o_1\}$
$CS_B = \{w_1o_1, w_5o_1\}$
$CG = \{q_1, r_1, p'_1\}$
$CS = \{w_1o_1, w_5o_1\}$

After (39), the stanceQUD is resolved and can be removed, because $CS \subset p'_1$. The participants have determined that they share an outlook that contains a positive stance towards p_1. Note nevertheless that they crucially have not determined whether p_1 is true.

We can also model non-canonical disagreements as in (40).

(40) QUD: *Which video is the newest?*
A: I think the first video is the newest.
B: I think the second video is the newest.

This conversation diverges from the one above after 37. B's assertion this time adds the commitment r_2 to DC_B, and A automatically follows suit.

(41) $\text{TABLE} = [\{r_1, r_2...\}, \{p'_1, p'_2...\}]$
$DC_A = \{q_1, r_2, p'_1\}$
$DC_B = \{q_1, r_2\}$
$CS_A = \{w_2o_1, w_6o_1\}$
$CS_B = \{w_2o_1, w_2o_2, w_6o_1, w_6o_2\}$
$CG = \{q_1, r_2\}$
$CS = \{w_2o_1, w_2o_2, w_6o_1, w_6o_2\}$

After (41), the subQUD is resolved and can be removed, because $CS \subset r_2$. The participants can also jointly infer from B's assertion that B also has a positive stance towards p_2. This positive stance p'_2 could be entered into B's discourse commitments, in turn restricting the outlooks the worlds that remain in their commitment set.

(42) $\text{TABLE} = [\{p'_1, p'_2...\}]$
$DC_A = \{q_1, r_2, p'_1\}$
$DC_B = \{q_1, r_2, p'_2\}$
$CS_A = \{w_2o_1, w_6o_1\}$
$CS_B = \{w_2o_2, w_6o_2\}$
$CG = \{q_1, r_2\}$
$CS = \{\}$

But (42) is catastrophic, with no possible joint outlooks remaining in CS. If participants remove inferred stance commitments, they could end up in the state in (43), no longer catastrophic but notably without any answer to the QUD on the table.

(43) $\text{TABLE} = [\{p'_1, p'_2...\}]$
$DC_A = \{q_1, r_2\}$
$DC_B = \{q_1, r_2\}$
$CS_A = \{w_2o_1, w_2o_2, w_6o_1, w_6o_2\}$
$CS_B = \{w_2o_1, w_2o_2, w_6o_1, w_6o_2\}$
$CG = \{q_1, r_2\}$
$CS = \{w_2o_1, w_2o_2w_6o_1, w_6o_2\}$

Because the participants still have pressure to establish a joint stance, and because they are free to make new attitudinal claims for times beyond t and t', a likely continuation is to attempt to convince someone to switch attitudes, and thereby establish a joint stance. For instance, B may eventually commit to some new proposition s_1, that B thinks p_1 at time t'', thereby offering a chance to infer that p'_1 should be added to their discourse commitments. This would result in (44), a late but successful resolution.

(44) $\text{TABLE} = [\{p'_1, p'_2...\}]$
$DC_A = \{q_1, r_2, s_1, p'_1\}$
$DC_B = \{q_1, r_2, s_1, p'_1\}$
$CS_A = \{w_1o_1, w_5o_1\}$
$CS_B = \{w_1o_1, w_5o_1\}$
$CG = \{q_1, r_2, s_1, p'_1\}$
$CS = \{w_1o_1, w_5o_1\}$