Shifting Opinions: An Experiment on Agreement and Disagreement in Dialogue

Shauna Concannon **Patrick G. T. Healey** **Matthew Purver**
Queen Mary University of London
Cognitive Science Research Group
School of Electronic Engineering and Computer Science
s.concannon@qmul.ac.uk

Abstract

Disagreement is understood to be socially problematic; it also rarely surfaces in naturally occurring conversation. An experiment was designed to allow us to directly manipulate the occurrence of exposed (dis)agreement and track its effects on the subsequent dialogue. This is the first experiment to directly manipulate the occurrence of exposed agreement and disagreement in dialogue. Insertions of exposed disagreement disrupt dialogues, bringing the topic of disagreement directly into the conversation, provoking clarification requests and resulting in a greater number of self-edits when formulating turns. The insertion of disagreement also led to more instances of exposed agreement, suggesting that dialogue partners co-operate to redress the face-threat of disagreement. Conversely, exposed agreement insertions were not as incongruous and had less disruptive impact on the ensuing dialogues; however, introducing agreement into the dialogue did lead to greater deliberation, with more alternative scenarios considered by participants during the task.

1 Introduction

Disagreeing or expressing a view in opposition to that of your interlocutor can be socially problematic. Disagreement has been associated with confrontation and conflict. Brown and Levinson (1987), in their seminal work on politeness, explain the predisposition for the avoidance of disagreement in terms of *face*, the concept derived from Goffman, relating to the public self-image or identity of an individual in interaction with others (Goffman, 1967). Direct challenges to a speaker or disagreeing with their assertion in dialogue can constitute, in Brown and Levinson's terminology, what is known as a *Face Threatening Act*, that is to say it can threaten the hearer's public self-image.

Conversation Analysts have shown that when people produce assessments of situations or events, positive responses are made more quickly and clearly than negative or unaligned responses (Sacks, 1987; Pomerantz, 1977). Negative or *dispreferred* responses are normally produced more slowly and are often prefaced with some form of agreement (e.g. 'Oh yes... but') and the negative assessment itself is often delayed by several turns and produced with some sort of mitigating account (Pomerantz, 1977). Disagreement, especially when done in a direct manner, is rare in conversation (Concannon et al., 2015). This means that it is difficult to assess what effects it has upon a dialogue. An experimental approach has the advantage that it allows us to directly manipulate the occurrence of exposed (dis)agreement and track its effects on the subsequent dialogue.

Previous studies on disagreement take a distributional or corpus based approach at evidencing and analysing instances of disagreement in interaction (Walker et al., 2012; Abbott et al., 2011; Misra and Walker, 2013; Holtgraves, 1997). These studies have provided valuable insights into the ways in which these complex social interactions are handled in different contexts, and given rise to various theories on how we process, respond to and mitigate the impact of disagreement. However, the literature also highlights that exposed disagreement rarely surfaces in naturally occurring conversation (Pomerantz, 1977; Concannon et al., 2015).

This paper outlines an experimental approach for investigating disagreement, which provides opportunity to manipulate the occurrence of *exposed* (dis)agreement in dialogue. By *exposed*, we refer particularly to direct and unequivocal presen-

tations of agreement and disagreement, such as 'I agree' and 'You're wrong'. However, we also explore less direct markers, which can, but do not always function in a (dis)agreement capacity. For example, turn initial 'no' and 'yes', can and are often used to signal agreement and disagreement, however, the function of these markers is context specific and depends on the preceding content (for example a 'no' following a negative statement can function as agreement).

1.1 Politeness and Accommodation Theory

One argument for the scarcity of disagreement in dialogue is anchored to the concept of politeness. Politeness Theory builds upon Ervin Goffman's concept of *face*. Goffman (1967) defines face as 'the positive social value a person effectively claims for himself' through interaction and offers a model of co-operation that is enacted when an individual's face or social value is threatened during interaction. Goffman stresses the co-operative nature of facework: 'When a face has been threatened ... lack of effort on the part of one person induces compensative effort from others' (Goffman, 1967). This mutual co-operation and shared consideration in interaction has also been located as a central notion for Politeness Theorists (Brown and Levinson, 1987; Watts, 2003).

Politeness Theory suggests that interlocutors minimise disagreement to save face, employing strategic conflict avoidance techniques to mitigate the effect of any disagreement that may surface (Leech, 1980). However, Accommodation Theory would posit that if someone is agreeable their conversational partner would match them in this convivial approach, whereas if they are adopting a discursive or even combative linguistic style, then their conversational partner would be likely to adopt a similar tact and synchronicity would become more exaggerated (Giles and Smith, 1979). Accommodation Theory posits that interlocutors adopt strategies of *convergence* to integrate and identify socially with another (Giles et al., 1991); this involves the adoption of linguistic similarities and leads to perceived communicative effectiveness (Giles and Smith, 1979) and cooperativeness (Feldman, 1968). Conversely, speech *divergence* reflects distancing from the co-conversant and can surface when confronted with perceived differences to the co-conversant.

1.2 Polite disagreement: When the context is right

Recent literature on disagreement and politeness theory in Sociolinguistics and Conversation Analysis suggests that in certain contexts disagreement is appropriate (Kotthoff, 1993), can signal sociability and intimacy (Schiffrin, 1984; Tannen, 1984; Angouri and Tseliga, 2010), and rather than lead to conflict, help strengthen relationships (Georgakopoulou, 2001; Sifianou, 2012). Furthermore, Chiu (2008) found in problem solving dialogues that disagreement, *when done politely*, was more productive in provoking novel contributions from participants than agreement. So although disagreement, particularly when executed impolitely, tends to be problematic, for certain contexts, such as problem solving and discussion tasks, it may be essential in advancing the deliberative quality of a dialogue. Chiu (2008) also suggests that *agreement* can be potentially detrimental to a dialogue, but the problematic aspects of agreement are not well reported in the literature; this gives rise to the question, 'what effect does exposed *agreement* have upon a dialogue?' If it is problematic, how and in what ways does this manifest?

If disagreement encourages novel contributions does agreement, conversely, stifle them? If people are too readily agreeing, does this prevent more involved discussion that could lead to shifts in stance or the development of new contributions? In order to understand the effects of both exposed agreement and disagreement, an experiment was designed that enabled the manipulation of such features under controlled conditions.

1.2.1 Can disagreement lead to more considered discussion?

A motivating factor behind this research is an interest in how individuals are led to shifts in stance, and how and when this occurs through interaction. Although there is good reason to think that disagreement ought to be socially problematic, as well as the insights provided by Chiu (2008), research on the phenomenon of repair shows that disruption in interaction can also be potentially beneficial to the progression of a dialogue (Healey, 2008; Colman et al., 2011), particularly if focused on the clarification of a content issue. Although instances of repair seemingly interrupt the flow of a dialogue, this attempt to address problem-

atic talk is not necessarily negative, rather it seems to drive the conversation forward. Issuing only agreements can often lead to a lack of mutual intelligibility in fact, which is why instances of repair are so common in task-oriented dialogues (Colman et al., 2011), a context where effective coordination is critical to the interactional outcome. Healey (2008) demonstrates that repair processes deal directly with misalignments and have a positive effect on measures of interactional outcome. Consequently, disagreement ought to be a catalyst or precursor to a potential shift in stance, as it signals a direct challenge to a held idea, which in turn may be retained, re-negotiated or or more fundamentally re-conceived. This, together with the findings by Chiu (2008), suggests that disagreement can play an important role in the deliberation and problem solving process.

1.3 Predictions

Given the literature we would expect that exposed disagreement would be especially problematic; it should instigate additional work being done in the interaction and more instances of repair. Insertions of exposed disagreement should be more disruptive than exposed agreement insertions, which should in turn facilitate more agreement. Assuming speakers are being co-operative, all things being equal, then disagreement should lead to more hedging and mitigation in order to manage the disagreement and minimise face threat. However, it may also lead to additional shifts in stance, or the consideration of more alternatives during the discussion dialogues.

2 Agreement and Disagreement Fragment Experiment

In order to assess the impact of exposed (dis)agreement, an experiment was designed in which instances of exposed (dis)agreement were artificially inserted into a dialogue. Turn-initial discourse markers such as 'No', 'But' 'You're wrong' and, 'I disagree' can highlight instances of disagreement within a conversation. Similarly, 'Yes', 'And', 'I agree' and 'You're right' can serve as indicators of agreement, or reinforce congruence. These eight fragments were selected because they provide a range of exposed, direct (dis)agreement and more subtle markers that *can* be used in (dis)agreement.

2.1 Hypotheses

1. Accommodation Theory: The general accommodation hypothesis is that dialogue partners match linguistic and discursive style. Thus the general *accommodation hypothesis* predicts that the insertion of agreement fragments will elicit additional instances of agreement, while the insertion of disagreement fragments will elicit additional instances of disagreement.

2. Politeness: The general politeness hypothesis is that face-threatening acts are socially problematic and should result in compensatory action being taken to redress and mitigate the situation. The general *politeness hypothesis* thus predicts that inserting disagreement fragments into a dialogue should lead to more work being done and more cooperation and consideration being displayed; this may result in increased effort when formulating responses (higher number of self-edits) and more clarification requests, expressions of agreement and other routinised polite sequences.

3. Productive Disagreements: The general productive disagreement hypothesis is that disagreement is essential for advancing the deliberative quality and problem solving aspects of dialogue. The *productive disagreement hypothesis* thus predicts that people will respond constructively to disagreement. The specific predictions for particular response measures are a much lower level issue, but we would expect the insertion of disagreement fragments to lead to increased deliberation taking place which lead to a higher number of shifts in stance over the course of a dialogue.

2.2 Method

Pairs of participants were seated at separate computers in adjacent rooms and given an instruction sheet to read detailing the balloon task. Participants are presented with a fictional scenario in which an hot air balloon is losing altitude and about to crash. The only way for any of three passengers to survive is for one of them to jump to a certain death. The three passengers are: Dr. Nick Riviera, a cancer scientist, Mrs. Susie Derkins, a pregnant primary school teacher, and Mr. Tom

Derkins, the balloon pilot and Susies husband. Participants are told to take as much time as they need to read the summary of the situation and then discuss with their partners via a chat tool set up on the computer at which they are seated, and attempt to come to a conclusion over who should jump from the balloon. The advantages of this task are that it is effective at generating debates between subjects and involves articulations of agreement and disagreement as they attempt to come to a conclusion. There is also plenty of scope for deliberation and shifts in stance.

2.3 Participants

Seventy-two participants were recruited, 46 female and 26 male, with the majority being undergraduate and postgraduate students at the University of London. Participants were invited to attend with someone who they already knew. They were recruited in pairs to ensure that inter-pair participants were acquainted. For a couple of experiments if one participant didn't show up a stand in was recruited last minute, and in these exceptions, which are marked in the data, the pair were not previously acquainted with each other. Each participant was paid at a rate of £7.50 per hour for participating in the experiment, or if they were a Psychology student at Queen Mary University of London then they could receive course credits in lieu of payment.

2.4 Materials

The participants communicate via a specially programmed chat tool, similar to other instant messenger interfaces they may have used previously. The Dialogue Experimental Toolkit (DiET) chat tool is a text-based chat interface facilitating real time manipulations of the dialogue. It is possible to programme several different types of interventions using the chat tool: turns may be altered prior to transmission, turns may not be relayed, and additional turns may be added, (e.g. Healey et al. (2003), insertion of spoof clarification requests).

These manipulations occur as the dialogue progresses, thus making them minimally disruptive to the sequence of dialogue. The DiET chat tool is built in Java and consists of a server console and user interface. Participants are faced with a text box displaying the conversation history and a smaller text box into which they can type. Participants can type simultaneously and their message is relayed to their conversation partner by use of the ENTER key. The server time stamps and stores all key presses. All turns are passed to the server before being transmitted to the other participant, thus making it an intermediary between what the participants type and what they receive. Turns can be automatically altered, removed or inserted by the server before they are relayed.

2.5 Design

The experiment is conducted in pairs; there were 12 dyads for each condition. Pairs of participants were presented with a discussion task and instructed to discuss for 30 minutes and attempt to come to an agreement. Each pair of participants was assigned to a condition at random. There were three experimental conditions. Please note, what we gloss here as the *Agreement* and *Disagreement* conditions, are named as such because the inserted fragments in each condition *can* index disagreement, however, we recognise that the more indirect fragments do not consistently perform this function.

- **Control condition**: Participants are welcomed and briefed before being sat at their respective computers, which were situated in adjoining rooms. They receive their task instructions on a piece of paper and can start when they are ready. They are instructed to discuss the scenario and attempt to come to an agreement on who should jump from the balloon for 30 minutes. No interventions are performed by the server; participants receive the dialogue turns exactly as they were typed.

- **Agreement condition**: Initial procedure is exactly the same as the control condition. Participants receive the dialogue turns exactly as they were typed, except for every fourth turn when one of the following fragments is inserted position: *you're right, I agree, yes, and*.

- **Disagreement condition**: Initial procedure is exactly the same as the control condition. Participants receive the dialogue turns exactly as they were typed, except for every fourth turn when one of the following fragments inserted at turn-initial position: *you're wrong, I disagree, no, but*.

A small scale pilot study raised some design implications that were accordingly addressed. In the

Agreement and Disagreement conditions manipulations were carried out every fourth turn issued by each speaker as this was deemed an acceptable frequency for interventions without proving too disruptive to the conversation. No intervention was made if the turn consisted of only one word, or the turn started with the same text as featured in the insertion fragments. This was to avoid the production of particularly non-sensical turns such as *you're wrong I agree*. The fragments were cycled through in order but the exposed (dis)agreement fragments (*you're wrong/right, I (dis)agree*) appeared half as often due to their marked nature.

3 Results

Data was gathered both directly from the chat tool which logged various features such as typing time, number of self-edits, i.e. use of the backspace and delete key and temporal data, as well as the transcripts themselves, which were analysed for linguistic features and frequencies. Additionally the resulting transcripts were hand coded for clarification requests and *stance shifts*, explained in more detail below.

3.1 A note on terminology

Turn: For the purpose of this experiment, a turn constitutes the text relayed in a single message, meaning what is delineated by the ENTER key.

Intervention Turn (IT): The IT refers to the turn issued by a speaker which has had a Turn-initial intervention fragment inserted before the actual typed message.

Intervention Reply Turn (IRT): The IRT refers to the next turn issued by the speaker who receives the Intervention Turn. This is not always the next sequential turn after the IT, as the speaker whose turn contained the IT may issues another turn.

Clarification Requests: The transcripts were hand coded for Clarification Requests (CR), a form of repair in which speakers signal a need for further information, typically due to a lack of full comprehension of a previous utterance. This was done by a single annotator, blind, and all labelling indicating which condition a file belonged to was removed. CRs were hand labeled in the dataset, based on Purver et al. (2003) schema, example provided in Table 1 .

Stance shifts: The transcripts were hand coded for shifts in stance regarding who to throw off of the balloon, i.e when a participant changed their point of view over who to sacrifice or save. There were seven potential stance states that cover all the possible combinations of who to save and who to sacrifice[1]. This was done by a single annotator, blind, and all labelling indicating which condition a file belonged to was removed. A participant's stance was carried over to the next turn, unless it provided new information that contradicted the previous stance.

Turn 1:	P1	**you're wrong** or maybe we are just going by gender stereotypes.. the feminist in me is screaming	IT
Turn 2:	P1	haha	
Turn 3:	P2	what if thats the whole point	IRT
Turn 4:	P1	sorry what if....?	CR
Turn 5:	P1	susie jumped?	CR

Table 1: Example of Reply Turn labelling

3.2 Overview of dataset

Table 2 displays the descriptive data for the turn, word and character counts for each condition.

	Condition		
Avg.	Control	Agreement	Disagreement
Turns by Dyad	86.71	63.17	70.79
Words by Dyad	587.67	555.58	535.08
Char. by Dyad	2938	2797	2710
Words per turn	7.41	9.49	9.11

Table 2: Summary of average typed data per condition

Both intervention conditions result in fewer overall turns than in the Control condition, but this was particularly the case, and statistically significant, with the Agreement condition (positive and agreement insertions, such as *yes* and *I agree*). Although the Agreement condition features fewer

[1]The range of possible stances: 1. Undecided, 2. Save Susie but undecided on who should die, 3. Save Nick but undecided on who should die, 4. Save Tom but but undecided on who should die, 5. Sacrifice Susie (and therefore save the other two), 6. Sacrifice Nick, 7. Sacrifice Tom.

turns than the Control condition, there are more words per turn on average in the Agreement condition. A non-parametric Kruskal Wallis test confirms a significant overall effect of Condition on the turns typed in the dialogues ($H_{(2)}$ = 6.34, p<0.04).[2] Subsequent planned pairwise comparisons with the Dunns test showed a significant increase in the number of turns per dyad in the Control condition compared to the Agreement condition (p<0.05). There is an overall effect of condition on the distribution of average words per turn, as confirmed by a non-parametric, Kruskal Wallis test ($H_{(2)}$ = 6.55, p<0.04). Subsequent planned pairwise comparisons with the Dunns test showed a significant increase between Agreement and Control conditions (p<0.03).

3.3 Message construction

Condition	Typing Time	Self-edits
Control	11850	6.98
Agree	16210	6.97
Disagree	13484	7.51

Table 3: Table depicting averageTyping Time and number of Self-edits (delete key presses), per turn, per condition

Table 3 shows the average typing time in milliseconds and the number of self-edits per turn. Self-edits are represented by the number of times the delete key is pressed during turn construction. A non-parametric Kruskal Wallis test finds an omnibus effect of condition on the number of self-edits during turn construction ($H_{(2)}$ = 40.92, p<0.01), with planned pairwise comparison revealing significant difference between the Agreement and Disagreement conditions (p<0.01). An overall effect of condition on typing time is confirmed by a non-parametric Kruskal Wallis test ($H_{(2)}$ = 99.28, p<0.01), with planned pairwise comparison revealing significant difference between the Agreement and Control conditions (p<0.01).

3.4 Message content

The following tables highlight differences in the content of the dialogues, such as Clarification Requests and instances of exposed and potential disagreement.

[2]Throughout we use p<0.05 as our criterion level but report computed probabilities to two decimal places for completeness.

3.4.1 Clarification Requests

Condition	Total Number of CRs	Mean CRs per dyad
Control	10	0.42
Agreement	13	0.54
Disagreement	50	2.08

Table 4: No. of Clarification requests by Condition

Table 4 shows the number of Clarification Requests by condition. The Disagreement condition has a significantly higher number of Clarification Requests than Control condition and Agreement condition. A non-parametric Kruskal Wallis test confirms an overall effect of Condition on the number of Clarification Requests in the dialogues ($H_{(2)}$ = 12.03, p<0.01). Planned pairwise comparison showed a significant increase between Control and Disagree conditions (p<0.01) and Agree and Disagree (p<0.02).

3.4.2 Instances of exposed and potential (dis)agreement

Table 5 shows the frequencies of turn-initial exposed and potential (dis)agreement markers. The markers included here are the same ones that feature in the fragments that were artificially inserted during the experiment.

Turn-initial	Control condition	Agreement condition	Disagreement condition
Exposed (dis)agreement			
I agree	2	5	10
You're right	0	0	0
I disagree	0	0	3
You're wrong	0	0	3
Totals:	2	5	16
Yes	170	124	139
No	29	23	35
And	103	51	55
But	119	81	77

Table 5: Table providing frequency data of turn-initial content of messages relayed during experiment dialogues.

Exposed (dis)agreement is more frequent in the Disagreement condition. A non-parametric Kruskal Wallis test shows a significant omnibus effect of condition on turn-initial exposed (dis)agreement ($H_{(2)}$ = 9.74, p<0.01). Subsequent planned pairwise comparisons with the Dunns test showed a significant increase in the number of instances of exposed (dis)agreement in the Disagreement condition compared to the control condition (p<0.01).

3.5 Deliberation and shifts in stance

The experiment transcripts were also hand labeled for *stance shifts*, i.e. when a participant voices a departure from one held opinion to an alternative regarding who should jump from the balloon.

Condition	Total	Median	Mean	St. Dev.
Control	175	7.5	8.33	3.96
Agree	248	11	10.33	4.88
Disagree	175	6	7.29	4.31

Table 6: Total number of stance state changes and averages per participant

The total number of state changes and average per participant by condition are shown in Table 6. The median number of stance state changes per participant is significantly effected by condition ($\chi(2)$ = 6.91, p=0.03). A Median Test was conducted as the variance is not approximately equal across samples, being much larger for the agreement condition. This result suggests that the Disagreement condition tends to reduce the number of alternatives people will consider and the agreement condition tends to increase it.

There is no correlation between the length of the conversation (in turns) and the number of state changes (Kendals Tau = -0.007, p = 0.94), so the significance is not related to nor skewed by the fact that the Agreement condition contains longer dialogues, i.e. it is not just about how much participants talk.

4 Discussion

The turn-initial frequency data shows that exposed agreement *and* disagreement are more common the Disagreement condition. The is counter to our Accommodation hypothesis, which anticipated that agreement would lead to more agreement while disagreement would engender more disagreement. Although there are notably zero instances of exposed disagreement in the Agreement condition, the comparative frequency in the Disagreement condition did not confirm a significant effect of condition. Furthermore, a third of the instances of turn-initial exposed disagreement in the Disagreement condition are actually instances of repair, rather than disagreement. As shown in the following excerpt from an experiment transcript, the exposed disagreement is incongruent, jarring and provokes a repair sequence. The respondent quotes back the source of trouble, indicated by the asterisks in the example below, which were falsely counted as turn-initial disagreement. The artificial insertions are shown in square brackets:

Example 1
A: Pros of keeping the doctor alive
A: [you're wrong] cures cancer
B: [no] you're wrong?*
A: What about?
B: no, I don't understand what you just said
B: You're wrong cures cancer?*
A: The doctor, if still alive will be about to discover the sure for the 'most common types of cancer'

This example demonstrates the disruptive nature of the inserted disagreement fragment; it disrupts the dialogue and is deemed incongruous enough for participant B to comment on, while participant A simply carries on with the conversation. This occurred several times in the dataset, however, only ever with the exposed disagreement fragments and never with the exposed agreement fragments. In line with the literature we found that exposed disagreement is especially problematic and on one occasion the insertion was so problematic that it was directly referenced and quoted by a participant, with both participants being alerted to the intervention.

Example 2
A: imagine how many scientists in the world
B: you're wrong theres a lot
A: i'm wrong?
B: what?
A: you said 'you're wrong theres a lot'
A: [no] what am i wrong about?

The Disagreement condition featured a significantly higher number of clarification requests. The Productive Disagreement hypothesis anticipated that this would signal additional work being done by participants trying to more fully understand one another's point of view. However, it is possible that the clarification requests are more clausal clarification than an attempt to understand the content; this interpretation is supported by the Example 2, which notably features a high number of clarification requests in a very short segment of dialogue, however, further analysis is needed to confirm this.

The Productive Disagreement hypothesis also anticipated that the Disagreement condition would lead to more *stance states* being considered. The results show that although there is an effect of condition on the number of different stance states or scenarios considered during the dialogue, the directionality was contrary to our predictions. The insertion of agreement fragments led to more shifts in stance. This may be due to the particularly marked and direct nature of the exposed disagreement fragments, which may have closed the discussion down. This would align with the CA and Politeness Theory literature, as well as Chiu (2008), which specifies that while *polite* disagreement may yield more novel contributions, impolite disagreement is always problematic.

Overall, our results most strongly confirm the Politeness hypothesis. Insertions of exposed disagreement had a disruptive effect upon the dialogues, producing confusion and clarification requests due to their unexpected and incongruous nature. Conversely, exposed agreement, even though also inserted randomly, did not disrupt the dialogue in the same way and were never explicitly addressed by a participant. The Disagreement condition produced significantly more instances of exposed agreement, which is most easily interpreted in terms of politeness, face and redressive action; with additional exposed disagreement being introduced into the dialogues, it seems that participants respond with cooperation and attempt to redress the potential affronts to face posed by the inserted fragments. As predicted there were more self-edits in the Disagreement condition, suggesting that participants were having to work harder to respond to the potentially face threatening insertions. Our results most strongly support the Politeness hypothesis and confirm that exposed disagreement is problematic and disruptive in dialogue.

Acknowledgments

This work was funded by EPSRC through the Media and Arts Technology Programme, an RCUK Doctoral Training Centre EP/G03723X/1. The authors would like to thank the reviewers for their helpful comments.

References

Rob Abbott, Marilyn Walker, Pranav Anand, Jean E Fox Tree, Robeson Bowmani, and Joseph King. 2011. How can you say such things?!?: Recognizing disagreement in informal political argument. In *Proceedings of the Workshop on Languages in Social Media*, pages 2–11. Association for Computational Linguistics.

Jo Angouri and Theodora Tseliga. 2010. "you have no idea what you are talking about!" from e-disagreement to e-impoliteness in two online fora. *Journal of Politeness Research. Language, Behaviour, Culture*, 6(1):57–82.

P. Brown and S.C. Levinson. 1987. *Politeness: Some Universals in Language Usage*. Studies in Interactional Sociolinguistics. Cambridge University Press.

Ming Ming Chiu. 2008. Flowing toward correct contributions during group problem solving: A statistical discourse analysis. *The Journal of the Learning Sciences*, 17(3):415–463.

Marcus Colman, Patrick GT Healey, et al. 2011. The distribution of repair in dialogue. In *Proceedings of the 33rd annual meeting of the cognitive science society*, pages 1563–1568.

Shauna Concannon, Patrick GT Healey, and Matthew Purver. 2015. Taking a stance: a corpus study of reported speech. In *Proceedings of the 19th SemDial Workshop on the Semantics and Pragmatics of Dialogue (GoDial)*, Semdial Workshop, page In Press.

Roy E Feldman. 1968. Response to compatriot and foreigner who seek assistance. *Journal of Personality and Social Psychology*, 10(3):202.

Alexandra Georgakopoulou. 2001. Arguing about the future: On indirect disagreements in conversations. *Journal of Pragmatics*, 33(12):1881–1900.

Howard Giles and Philip Smith. 1979. *Accommodation theory: Optimal levels of convergence*.

Howard Giles, Justine Coupland, and Nikolas Coupland. 1991. *Contexts of accommodation: Developments in applied sociolinguistics*. Cambridge University Press.

Erving Goffman. 1967. On face-work. *Interaction ritual*, pages 5–45.

Patrick Healey. 2008. Interactive misalignment: The role of repair in the development of group sublanguages. *Language in Flux. College Publications*, 212.

Thomas Holtgraves. 1997. Yes, but... positive politeness in conversation arguments. *Journal of Language and Social Psychology*, 16(2):222–239.

Helga Kotthoff. 1993. Disagreement and concession in disputes: On the context sensitivity of preference structures. *Language in Society*, 22:193–193.

Geoffrey N Leech. 1980. *Explorations in semantics and pragmatics*. John Benjamins Publishing.

Amita Misra and Marilyn A Walker. 2013. Topic independent identification of agreement and disagreement in social media dialogue. In *Conference of the Special Interest Group on Discourse and Dialogue*, page 920.

Anita Pomerantz. 1977. *Agreeing and disagreeing with assessments: Some features of preferred/dispreferred turn shapes*. Centre for Socio-Legal Studies.

Matthew Purver, Jonathan Ginzburg, and Patrick Healey. 2003. On the means for clarification in dialogue. In *Current and new directions in discourse and dialogue*, pages 235–255. Springer.

Harvey Sacks. 1987. On the preferences for agreement and contiguity in sequences in conversation. *Talk and social organization*, 54:69.

Deborah Schiffrin. 1984. Jewish argument as sociability. *Language in society*, 13(03):311–335.

Maria Sifianou. 2012. Disagreements, face and politeness. *Journal of Pragmatics*, 44(12):1554–1564.

Deborah Tannen. 1984. The pragmatics of cross-cultural communication. *Applied Linguistics*, 5(3):189–195.

Marilyn A Walker, Jean E Fox Tree, Pranav Anand, Rob Abbott, and Joseph King. 2012. A corpus for research on deliberation and debate. In *LREC*, pages 812–817.

Richard J Watts. 2003. *Politeness*. Cambridge University Press.