COGNITIVE SCIENCE
A Multidisciplinary Journal

Cognitive Science 41 (2017) 259–271

ISSN: 0364-0213 print / 1551-6709 online
DOI: 10.1111/cogs.12325

Language Use and Coalition Formation in Multiparty Negotiations

Eyal Sagi,[a] Daniel Diermeier[b]

[a]*Psychology Department, University of St. Francis*
[b]*Harris School of Public Policy, University of Chicago*

Received 6 March 2014; received in revised form 19 May 2015; accepted 15 September 2015

Abstract

The alignment of bargaining positions is crucial to a successful negotiation. Prior research has shown that similarity in language use is indicative of the conceptual alignment of interlocutors. We use latent semantic analysis to explore how the similarity of language use between negotiating parties develops over the course of a three-party negotiation. Results show that parties that reach an agreement show a gradual increase in language similarity over the course of the negotiation. Furthermore, reaching the most financially efficient outcome is dependent on similarity in language use between the parties that have the most to gain from such an outcome.

Keywords: Negotiation; Coalition formation; Linguistic entrainment; Psycholinguistics; Latent semantic analysis

1. Introduction

Negotiations are usually conducted through the use of language. For negotiations to be successful the negotiating parties need to build common ground and represent the issues of the negotiation similarly. Pickering and Garrod (2004) suggest that such a convergence in representation is often the result of successful dialog accompanied by a convergence in patterns of language use. Language similarity, however, can also be the consequence of strategic behavior. In this article, we examine how similarity in language use evolves throughout the course of three-party negotiations and compare the language used by parties that reach agreement and those that are excluded from it.

Negotiations between multiple participants are more complex than two-party negotiations (Bazerman, Curhan, Moore, & Valley, 2000). This is especially true when

Correspondence should be sent to Eyal Sagi, Psychology Department, University of St. Francis, 500 Wilcox St., Joliet, IL 60435. E-mail: esagi@stfrancis.edu

an agreement requires only a subset of the negotiating parties, in which case being excluded from an agreement is an ongoing concern for each negotiator. Partial coalition agreements, however, are often less desirable than agreements that involve the group if they do not allocate all of the available resources. However, even being part of a partial agreement is preferable to not reaching an agreement or being excluded from an agreement reached by others (Diermeier, Swaab, Medvec, & Kern, 2008).

The added complexity of multiparty negotiation has been shown to affect the patterns of language use in such negotiations. Following the framework of communication accommodation theory (cf. Giles, Willemyns, Gallois, & Anderson, 2007), Huffaker, Swaab, and Diermeier (2011) demonstrate how the formation of a coalition is affected by specific aspects of the language used by the negotiating parties. Specifically, they find that partners to a coalition show more similarity in language use than participants who were not part of a coalition. The use of assents was also found to correlate positively with participating in the coalition agreement. In contrast, the use of negative emotion words was a detrimental predictor to being a part of a coalition. These results are congruent with empirical findings in psycholinguistics which show that in successful dialogs the representations and language used by dialog partners tend to converge over time (e.g., Brennan & Clark, 1996; Pickering & Garrod, 2004). Even intentionally created similarity in language use can facilitate agreement. For example, Swaab, Maddux, and Sinaceur (2011) asked participants to mimic the distinctive linguistic patterns of their negotiation partners. In particular, participants were asked to use the same jargon, abbreviation, and emoticons used by the other participant in the negotiation. Swaab et al. (2011) found that such intentional mimicry was effective when used early in the negotiation, but not in later stages.

Consequently, there are two possible patterns of language similarity that might lead to an eventual coalition: First, if strategic mimicry is an effective tool for reaching agreement, we might expect that early similarity in language use might lead to the forming of a coalition later on. Second, if similarity in language use is a consequence of a gradual alignment of representations and the emergence of a mutual understanding, then language similarity should increase over the course of the negotiation and ultimately result in the parties reaching an agreement.

It is also possible that both of these factors play a role. In that case, we would expect to find not only that eventual coalition partners show more similar language early on than non-coalition partners, but also that this difference increases over time.

While Huffaker et al. (2011) show that similarity in language use is correlated with the outcome of the negotiation, they use the entire negotiation as their unit of analysis. Consequently, their results do not explore whether such similarity exists early in the negotiation or whether it develops throughout its course. In this article, we measure language similarity between adjacent dialog moves, a measure that allows us to track changes in language similarity over the course of the negotiation.

As mentioned earlier, reaching an optimal agreement in a multiparty negotiation often requires the formation of a grand coalition, a coalition that includes all of the parties. However, reaching this kind of agreement can be difficult and may require a high level

of cooperation between the parties (Diermeier et al., 2008). Consequently, we are interested not only in when the parties use similar language, but also which parties' exhibit similarity in language use when more than two of them participate in the final coalition. Specifically, cooperation between participants increases their bargaining power and allows participants with weak bargaining positions, who might otherwise be excluded from the final coalition, to strengthen their position. Consequently, we hypothesize that participants in weaker bargaining positions have the most to gain from cooperating in order to achieve a grand coalition. Such a partial coalition should result in increased similarity in language use between these weaker parties.

1.1. Measuring similarity in language use

The measure of language similarity we use in this article is based on the latent semantic analysis (LSA) cosine similarity of a pair of utterances. Such a measure has been used in the past as a measure of textual coherence (Foltz, Kintsch, & Landauer, 1998; McNamara, Cai, & Louwerse, 2007) and as a measure of linguistic entrainment (Huffaker, Jorgensen, Iacobelli, Tepper, & Cassell, 2006). Measures based on LSA have also been used to explore a diverse set of psychological phenomena related to texts and language, such as semantic priming (Chwilla & Kolk, 2002; Landauer & Dumais, 1997), the acquisition of knowledge from texts (Wolfe et al., 1998), the representation of moral concerns (Sagi & Dehghani, 2014), and conceptual framing (Sagi, Diermeier, & Kaufmann, 2013).

Latent semantic analysis vectors for individual words are generated based on the co-occurrence patterns of words in large corpora. These vectors identify points in a high-dimensional space.[1] The more likely two words are to co-occur with similar words, the closer they will be in the space. For example, the vectors for *sun* and *moon* are fairly close together and show a cosine similarity of .53, whereas *man* and *moon* are not very similar and show a cosine similarity of .03. Moreover, when several word vectors from a single utterance are combined together, as was done in this study, the result identifies a point in space that represents the overall topic of the utterance.

It is important to note that this kind of automatic measure disregards certain linguistic elements that a human coder might use. For instance, the use of negation is generally ignored, while sarcasm and metaphors are often misrepresented. However, since we are interested in the convergence of language use—that is, whether participants are using similar and/or related words and terms to convey their (sometimes conflicting) ideas, this type of analysis seems sufficient.

2. Method

2.1. The corpus

The data used in this article come from a study reported by Huffaker et al. (2011). They patterned their study after a coalition game developed by Raiffa (1982). In that

study, 180 MBA students were divided into 60 three-person groups. The participants in the study had some experience in managerial decision making and negotiation. Within each group, participants were assigned to one of three roles (A, B, C) and instructed that they were to use an online chat room to negotiate a split of that payoff among themselves. Participants were unaware of the identities of the other participants in the negotiation.

All participants were provided with the payoff table in advance of the negotiation (see Table 1) and informed of their role in the negotiation and the role of the other participants. As is evident from the table, different coalition formations received different payoffs, and if no coalition was formed all participants received a payoff of zero. The participants were allowed to negotiate how the payoff was distributed between them. These payoff options provided incentives for the participants to join up with another participant in a partial coalition and then jointly take advantage of the resulting weak bargaining position of the third participant. However, the payoff table was designed so that the third player left out of a partial coalition could always make an attractive offer to one of the members of the initial coalition to induce a defection from the preliminary agreement. Consequently, participants were incentivized not only to be a part of a forming coalition, but also to ensure that it is a stable coalition and that their partner(s) would not defect. Only then were they able to use their bargaining power to extract a favorable outcome from the remaining third party. If members of a two-party coalition were unsure about the reliability of their respective coalition partner, they may have refrained from engaging discussions with the third party for fear of giving the "out party" an opening to split the current coalition. In this case, the negotiation resulted in a partial coalition leaving some money on the table, compared to the most efficient coalition, which allocated the largest monetary award among all parties. Consequently, the ideal strategy for each participant was to quickly form a stable two-party coalition and squeeze the remaining party into accepting an unfavorable deal. Maintaining such a position of strength, however, required maintaining the loyalty of the coalition partner in the process.

Each session of the experiment lasted an hour. Participants in the experiment were placed at computers in different rooms so that their only means of communication with

Table 1
Payoff table in the negotiation game from Huffaker et al. (2011)

Possible Agreements	Total Payoff
A alone	$0
B alone	$0
C alone	$0
A and B	$118,000
A and C	$84,000
B and C	$50,000
A, B, and C	$121,000

Notes. A, B, and C represent the participants in the negotiation. The payoff is split between the parties that reach the described final agreement according to their agreement, not necessarily evenly.

each other was through the provided chat software. They logged into a public chat room to begin the negotiation process. The negotiation ended when at least two participants reached an agreement and declared it as final, or the allotted time had elapsed.

The software allowed participants to move from the public chat room to three private chat rooms. That is, participant A could move into one of the private chat rooms together with participant B so that they could negotiate without participant C being privy to the content of the negotiation. However, all participants were alerted whenever a participant entered or exited a chat room so that the excluded participant was always aware that the two other participants might be negotiating in private. This mimics some of the real-world aspects of a negotiation, where parties are often able to communicate in private, but the fact that they communicate in private is common knowledge. A private exchange of information can also provide an indication that the two parties are forming a coalition.

While it is theoretically possible for the two parties to exclude a third party completely from the negotiation process, such exclusionary behavior rarely occurred. In particular, utterances in the public chat room were fairly evenly distributed across the quarters of the negotiation and there were generally more utterances in the public chat rooms than in the private rooms.

Sessions lasted 33–318 turns (M = 110.10, SD = 53.31). The most common coalition was AB (24 times). No agreement was reached in five sessions. Most sessions lasted about 40–60 min, while the shortest session lasted about 20 min and concluded with a grand coalition.

2.2. Semantic analysis

The analysis in this article is based on the transcripts of the negotiation experiment conducted by Huffaker et al. (2011). Each submitted line in the chat was considered a separate utterance. An LSA vector was computed for each individual utterance by using vector addition to combine the vectors of all of the content-bearing words in the utterance. When an utterance did not include any content-bearing words, a null vector was used to represent it. The vector space used for this analysis was generated by Infomap (Schütze, 1997; Takayama, Flournoy, Kaufmann, & Peters, 1998) using the written portion of the British National Corpus (BNC Consortium, 2007).[2] We then computed the correlation of vectors representing non-null adjacent utterances that occurred in the same chat context (i.e., when they occurred in the same chat room, whether public or private).[3] Fig. 1 presents some examples of utterance pairs and their correlations.

The first few utterances were of an introductory nature, mostly with participants reporting "I am present." These utterances consequently resulted in identical vectors for all three participants that inflated the correlations at the beginning of many sessions. Consequently, we elected to discard utterances that included the words "present," "hello," "hi," "morning," and "welcome."[4] There were 253 such utterances (3.9% of the corpus), all of which were greetings and occurred within the first 10% of each session. Moreover, 196 (77.5%) of these occurred within the first five utterances of a session and 124 (49%) are the formulaic utterance "I am present."

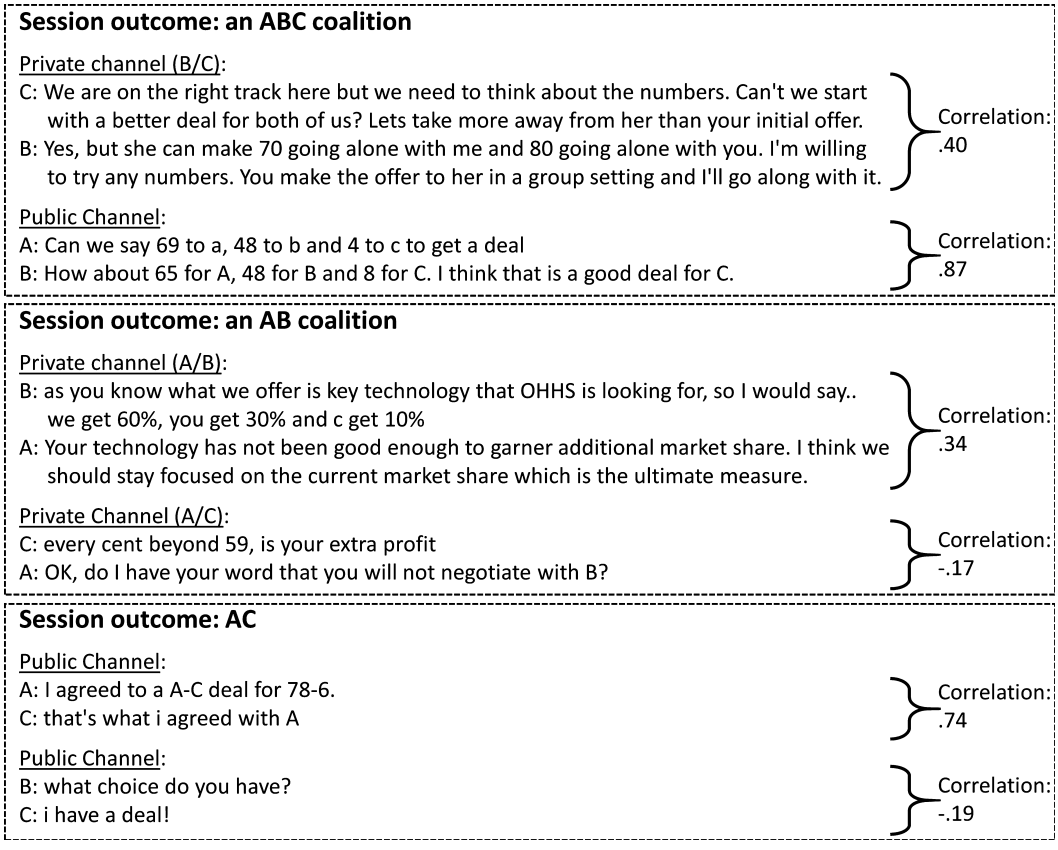

Fig. 1. Sample utterance pairs from three sessions and the computed correlations between them.

In order to measure language similarity, we categorized the utterance pairs based on the two participants that contributed to them. We predicted that participants who were included in the final coalition would use more similar language with each other than with participants who were excluded from the coalition. For example, if an AB coalition was ultimately reached, utterance pairs between A and B would be predicted to have more similar language use (i.e., utterance-to-utterance correlation) than those between A and C or B and C. Consequently, we divided the utterance pairs into those in which both participants were included in the final coalition (*intracoalition utterance pairs*) and those in which at least one of the participants was excluded from the coalition (*extracoalition utterance pairs*). Importantly, when the final coalition included all parties, all of the utterance pairs were considered to be intracoalition pairs. In contrast, when no agreement was reached, all of the utterances were considered to be extracoalition pairs.

We were also interested in examining how the difference in language similarity between intracoalition and extracoalition utterance pairs evolved over time. Therefore, we divided the utterance pairs based on their position in the negotiation. We used a relatively

coarse grain division of time (quarters) because some of the discussions consisted of relatively few utterances (under 50).[5] We categorized each utterance pair based on the quarter of the negotiation in which the first utterance of the pair occurred.

3. Results

We tested three distinct hypotheses:

1. Following accounts of linguistic entrainment (e.g., Pickering & Garrod, 2004), we hypothesized that coalition formation will be accompanied by the alignment of language use. On the basis of this hypothesis we predicted that as a coalition comes together, the language used by its participants will increase in similarity. Because forming a coalition takes time, we expected that this increase in similarity was more likely to occur later in the negotiation. Moreover, this increase in language similarity should not affect language use by excluded parties. This hypothesis therefore predicted that as the negotiation proceeds, the similarity of intracoalition utterance pairs, but not extracoalition ones, should increase.
2. Following the literature on the effectiveness of strategic mimicry in negotiations (e.g., Swaab et al., 2011), we hypothesized that a coalition might be more likely to form when one of the participants intentionally created linguistic similarity with another participant during the initial stages of the beginning of a negotiation. While this type of linguistic similarity would have been intentional, it might have led the second party to assume that it indicated a shared perspective and made a protocoalition between the two parties more likely. If this type of intentional manipulation is an effective tool in these negotiations, intracoalition utterance pairs should have been more similar than extracoalition ones early in the negotiation.
3. On the basis of Diermeier et al.'s (2008) theory of coalition formation, we hypothesized that a protocoalition involving participant A will be less likely to result in a grand coalition (ABC) than one that involves only participants B and C. In the negotiation game used by Huffaker et al. (2011), a protocoalition that involves participants B and C can increase its payoff by $71,000 by adding participant A to the final coalition. In contrast, protocoalitions that involve participant A have a substantially lower incentive to make the effort to incorporate the remaining participant because the increase in their payoff is substantially lower ($3,000 for an AB protocoalition, and $37,000 for an AC protocoalition). After all, protocoalitions have to balance the marginal payoff increase from adding the third player with the risk that the third player may try to break up the protocoalition. Consequently, a grand coalition is more likely to result from a BC protocoalition than either an AB or AC one. We therefore predicted that the pattern of language similarity for a grand coalition should be similar to that of a BC coalition. That is, a grand coalition should demonstrate a pattern where BC utterance pairs show a higher degree of similarity than AB or AC ones.

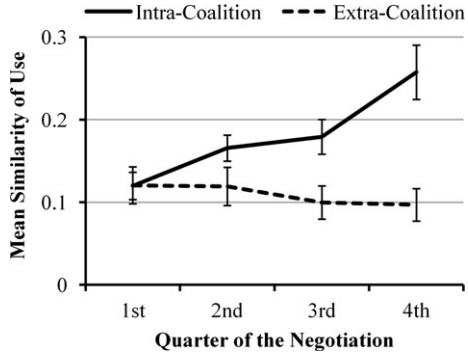

Fig. 2. Similarity of language use by utterance pair type and time for negotiation. Error bars represent standard error.

To test the first two hypotheses, we used a mixed model with the type of utterance pair (*intracoalition* vs. *extracoalition*) and its position in the session (first through fourth quarters) as the independent variables. The experimental session was included as a random variable. The dependent measure was the average utterance-to-utterance correlation. Fig. 2 plots the mean utterance similarity over time.

As expected, intracoalition utterance pairs (M = .19, SD = .17) showed more language similarity than extracoalition ones (M = .10, SD = .13; $F(1, 286) = 20.93$, $MSE = 0.02$, $p < .0001$). Overall utterance pair similarity (combining over utterance pairs from all parties) increased over time ($F(1, 286) = 9.21$, $MSE = 0.02$, $p < .01$). Most important, the difference in language similarity between intra- and extracoalition utterance pairs increased over time ($F(3, 286) = 13.71$, $MSE = 0.02$, $p < .001$). To further explore this interaction, Table 2 presents the results of planned comparisons of language similarity for intra- and extracoalition utterances for each quarter. Moreover, intracoalition utterance pairs showed an increase in language similarity over time ($F(3, 145) = 18.11$, $MSE = .024$, $p < .0001$), whereas no such increase was observed for extracoalition utterance pairs ($F(3, 102) < 1$, *n.s.*).

The gradual increase in language similarity over the course of the negotiation that was observed for intracoalition pairs is congruent with accounts in which a gradual alignment in language use and semantic representation is related to the likelihood of forming a coalition (Hypothesis 1). However, we found no evidence to support accounts in which early similarity in language use leads to the formation of a coalition (Hypothesis 2).

Table 2
Planned comparisons of similarity in language use between intra- and extracoalition utterance pairs

Quarter	*df*	*F*	*MSE*	*p*
1	1, 29	<1	.007	.6831
2	1, 26	3.99	.010	.0564
3	1, 25	7.28	.013	.0123
4	1, 28	14.66	.037	.0007

Huffaker et al. (2011) found that the use of assents, but not negations, was related to the formation of a coalition. Since such use might also affect linguistic similarity, we tested whether assents and negations affected similarity using a mixed model with the individual utterance-to-utterance similarity as the dependent measure. The independent measures were position, type of pair, and the use of assents and/or negations. We determined whether assents or negations were used by identifying whether words from the appropriate LIWC dimensions (Linguistic Inquiry and Word Count; Pennebaker, Chung, Ireland, Gonzales, & Booth, 2007) were present in the utterance pair.[6] The experimental session was included as a random variable. The similarity of language use also increased when assents were used ($F(1, 1570) = 21.16$, $MSE = 0.06$, $p < .0001$) and marginally decreased when negations were used ($F(1, 1570) = 2.85$, $MSE = 0.06$, $p = .09$). In the case of assents, this effect increased over time ($F(1, 1570) = 3.86$, $MSE = 0.06$, $p < .05$). There was no significant interaction between the type of pair and the use of either negations or assents. Most important, like the previous analysis, this model identified an interaction between the type of pair and the position of the utterance even after controlling for the effects of assents and negations ($F(1, 1570) = 11.28$, $MSE = 0.06$, $p < .001$).

We now turn to Hypothesis 3. To test this hypothesis, we compared the similarity in language use of the various pairs in sessions that resulted in a grand coalition to those that resulted in an AB coalition. We discarded sessions that resulted in other coalitions because there were not enough data to use them to conduct a meaningful analysis—there were only 10 sessions that resulted in an AC coalition and 4 that resulted in a BC coalition.

We used a mixed model with the participants of utterance pair (AB, AC, or BC), the utterance pair's position the negotiation session (first through fourth), and the resulting coalition (AB or ABC) as the independent variables. As stated earlier, the dependent measure was the average utterance-to-utterance correlation and the session was included as a random variable. Fig. 3 shows the mean utterance similarity for each participant pair for an AB coalition, and Fig. 4 shows the mean utterance similarity for each participant pair for an ABC coalition.

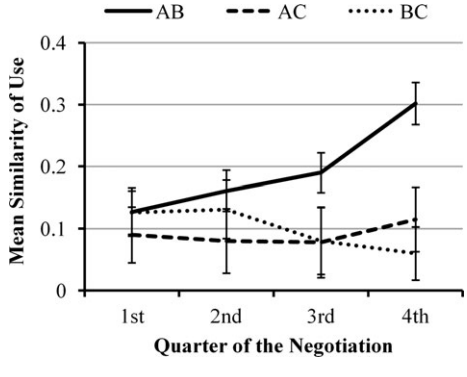

Fig. 3. Similarity of language use by utterance pair participants and time for negotiations that resulted in an AB coalition. Error bars represent standard error.

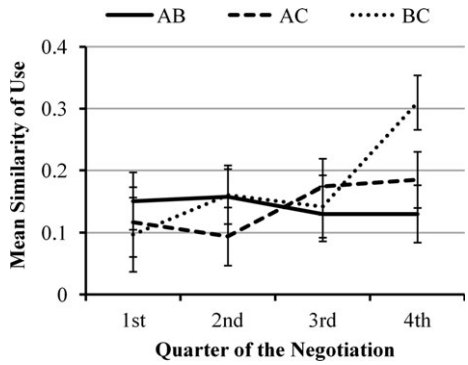

Fig. 4. Similarity of language use by utterance pair participants and time for negotiations that resulted in an ABC coalition. Error bars represent standard error.

The analysis yielded two statistically significant interactions—First, the three-way interaction between utterance pair participants, position within the negotiation, and the resulting coalition was significant ($F(2, 314) = 9.96$, $MSE = 0.03$, $p < .0001$). More important, the expected two-way interaction between the utterance pair participants and the resulting coalition was also significant ($F(2, 314) = 7.56$, $MSE = 0.03$, $p < .001$). Planned comparisons revealed that these interactions were most likely due to an increase in the language similarity between participants B and C during the later stages of negotiations that resulted in an ABC coalition ($F(1, 66) = 10.60$, $MSE = 0.03$, $p < .01$) and an increase in the language similarity between participants A and B in the later stages of negotiations that resulted in sessions that resulted in an AB coalition ($F(1, 106) = 8.56$, $MSE = 0.03$, $p < .01$). Moreover, the initial positions of the three possible pairs were not statistically distinguishable ($F < 1$ in both cases). No difference in language similarity between A and C was found between negotiations that resulted in ABC coalition and those that resulted in an AB coalition ($F(1, 67) = 1.10$, *n.s.*).

4. Discussion

The analysis presented here, based on data collected by Huffaker et al. (2011), demonstrates that as agreement forms during a negotiation, the language used by the parties involved becomes more similar. Moreover, parties that are excluded from the coalition do not show this pattern of convergence. This result has implications for theories of linguistic entrainment as it demonstrates that merely being party to a linguistic exchange is not enough. This suggests that linguistic similarity is representative of the alignment of representation rather than merely a result of exposure to linguistic input.

Our results are also congruent with Diermeier et al.'s (2008) model of coalition formation, in which building trust between parties plays a pivotal role, provided that language similarity, as a measure of shared perspective, is indicative of trust between the parties. Following this interpretation, trust appears to gradually emerge over the course

of a negotiation, and it is a reliable predictor of the ultimate outcome of the negotiation. Moreover, our results are consistent with an interpretation that trust between the participants with the most to gain from a grand coalition is the key to achieving such a coalition.

We found no evidence for the effectiveness of early similarity in language use in directing the course of a negotiation. However, this could be a consequence of the design of the negotiations we used. Research that demonstrates the effectiveness of early similarity in language use often relies on an explicit and strategic use of repetition and mimicry by the participants (e.g., Swaab et al., 2011). In contrast, the participants whose negotiations we analyzed in this article were not instructed to use such strategies. While strategic use of mimicry may be effective, we found no evidence that this strategy was used.

It is important to note that the negotiation task we analyzed in this article focuses on the split of a monetary payoff. While this task is commonly used to simulate negotiations in economics, many real-world negotiations involve various non-monetary stakes (e.g., work practices) resulting in a more complex negotiation task. Nevertheless, the fact that our results show evidence of linguistic alignment even in negotiations with minimal content suggests that the effects should be even more pronounced in negotiations that involve more diverse interests. Interestingly, as the statement from participant B in the AB coalition in Fig. 1 shows, some participants did make arguments that used non-monetary values (e.g., "key technology"), although they were infrequent.

Finally, our analysis illustrates that multiparty negotiations, while more complex than two-party negotiations and dialogs, follow many of the same patterns as their simpler counterparts. However, the added dynamics of such a negotiation allows researchers to examine topics, such as the role of similarity in language use in reaching agreement and the influence of the availability of private communication channels on the course of a negotiation, that are often difficult to explore when only two parties are involved in a linguistic exchange.

Acknowledgments

We thank David Huffaker and Roderick Swaab for sharing the data from their study.

Notes

1. In accordance with the accepted practices for Infomap (http://infomap-nlp.sourceforge.net/), we used 100 dimensions in our analysis.
2. This corpus was selected because it represents a well-balanced sample of the English language. Nevertheless, the choice of corpus does not make a large difference and a qualitatively and quantitatively similar pattern of results was obtained using a space derived from a corpus of New York Times articles (Sandhaus, 2008).

3. Because the first dimension of LSA vector spaces tends to correlate with the frequency and length of the text, it was dropped from the analysis (cf. Hu, Cai, Wiemer-Hastings, Graesser, & McNamara, 2007).
4. These are common and formulaic terms used at the beginning of the conversation to announce arrival. As such, they are generally used by all participants and do not constitute mimicry as defined by Swaab et al. (2011). Furthermore, including these terms inflates the correlations in the first part of the negotiations but does not otherwise change the results of the analysis presented below.
5. This coarse division is adequate for our hypotheses. It is possible to utilize smaller time units in an analysis of this type to gain further insight into the temporal progression of the negotiation. Essentially, the choice of temporal units for analysis represents a trade-off between precision and statistical power. Nevertheless, similar results can be obtained by analyzing individual utterance pairs based on their relative position in the session rather than aggregating them based on the quarter in which they appear.
6. LIWC is a tool that uses the frequency of occurrence of various words to estimate related linguistic and psychological variables, such as the use of assents and negations.

References

Bazerman, M. H., Curhan, J. R., Moore, D. A., & Valley, K. L. (2000). Negotiation. *Annual Review of Psychology*, *51*(1), 279–314. doi:10.1146/annurev.psych.51.1.279

BNC Consortium. (2007). *The British National Corpus*, version 3 (BNC XML Edition). Oxford University Computing Services. Available at http://www.natcorp.ox.ac.uk/. Accessed May 19, 2015.

Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*(6), 1482–1493. doi:10.1037/0278-7393.22.6.1482

Chwilla, D. J., & Kolk, H. H. J. (2002). Three-step priming in lexical decision. *Memory & Cognition*, *30*(2), 217–225. doi:10.3758/BF03195282

Diermeier, D., Swaab, R. I., Medvec, V. H., & Kern, M. C. (2008). The micro-dynamics of coalition formation. *Political Research Quarterly*, *61*(3), 484–501. doi:10.1177/1065912908316981

Foltz, P. W., Kintsch, W., & Landauer, T. K. (1998). The measurement of textual coherence with latent semantic analysis. *Discourse Processes*, *25*(2–3), 285–307. doi:10.1080/01638539809545029

Giles, H., Willemyns, M., Gallois, C., & Anderson, M. C. (2007). Accommodating a new frontier: The context of law enforcement. In K. Fiedler (Ed.), *Social communication* (pp. 129–162). New York: Psychology Press.

Hu, X., Cai, Z., Wiemer-Hastings, P., Graesser, A. C., & McNamara, D. S. (2007). Strengths, limitations, and extensions of LSA. In T. K. Landauer, D. S. McNamara, S. Dennis, & W. Kintsch (Eds.), *Handbook of latent semantic analysis* (pp. 401–426). Mahwah, NJ: Erlbaum.

Huffaker, D., Jorgensen, J., Iacobelli, F., Tepper, P., & Cassell, J. (2006). Computational measures for language similarity across time in online communities. In E. Hovy, K. Zechner, & L. Zhou (Eds.), *Proceedings of the HLT-NAACL 2006 workshop on analyzing conversations in text and speech* (pp. 15–22). New York: Association for Computational Linguistics.

Huffaker, D. A., Swaab, R., & Diermeier, D. (2011). The language of coalition formation in online multiparty negotiations. *Journal of Language and Social Psychology*, *30*(1), 66–81. doi:10.1177/0261927X10387102

Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*(2), 211–240. doi:10.1037/0033-295X.104.2.211

McNamara, D. S., Cai, Z., & Louwerse, M. M. (2007). Optimizing LSA measures of cohesion. In T. K. Landauer, D. S. McNamara, S. Dennis, & W. Kintsch (Eds.), *Handbook of latent semantic analysis* (pp. 379–400). Mahwah, NJ: Erlbaum.

Pennebaker, J., Chung, C., Ireland, M., Gonzales, A., & Booth, R. (2007). *The development and psychometric properties of LIWC2007*. Austin, TX: LIWC. Available at www.liwc.net. Accessed May 19, 2015.

Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, *27*(2), 169–189.

Raiffa, H. (1982). *The art and science of negotiation*. Cambridge, MA: Harvard University Press.

Sagi, E., & Dehghani, M. (2014). Measuring moral rhetoric in text. *Social Science Computer Review*, *32*(2), 132–144. doi:10.1177/0894439313506837

Sagi, E., Diermeier, D., & Kaufmann, S. (2013). Identifying Issue Frames in Text. *PLoS ONE*, *8*(7), e69185.

Sandhaus, E. (2008). *The New York Times annotated corpus*. Philadelphia, PA: Linguistic Data Consortium.

Schütze, H. (1997). *Ambiguity resolution in language learning: Computational and cognitive models*. Stanford, CA: CSLI Publications.

Swaab, R. I., Maddux, W. W., & Sinaceur, M. (2011). Early words that work: When and how virtual linguistic mimicry facilitates negotiation outcomes. *Journal of Experimental Social Psychology*, *47*(3), 616–621. doi:10.1016/j.jesp.2011.01.005

Takayama, Y., Flournoy, R., Kaufmann, S., & Peters, S. (1998). *Information mapping: Concept-based information retrieval based on word associations*. Stanford, CA: CSLI Publications.

Wolfe, M. B. W., Schreiner, M. E., Rehder, B., Laham, D., Foltz, P. W., Kintsch, W., & Landauer, T. K. (1998). Learning from text: Matching readers and texts by latent semantic analysis. *Discourse Processes*, *25*(2–3), 309–336. doi:10.1080/01638539809545030