University of Groningen

In the Short Term We Divide, in the Long Term We Unite

Maes, Michael; Flache, Andreas; Takacs, Karoly; Jehn, Karen A.

Published in:
Organization Science

DOI:
[10.1287/orsc.1120.0767](https://doi.org/10.1287/orsc.1120.0767)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2013

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Maes, M., Flache, A., Takacs, K., & Jehn, K. A. (2013). In the Short Term We Divide, in the Long Term We Unite: Demographic Crisscrossing and the Effects of Fault lines on Subgroup Polarization. *Organization Science*, 24(3), 716-736. <https://doi.org/10.1287/orsc.1120.0767>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

In the Short Term We Divide, in the Long Term We Unite: Demographic Crisscrossing and the Effects of Faultlines on Subgroup Polarization

Michael Mäs

Chair of Sociology, in particular of Modeling and Simulation, ETH Zurich, 8092 Zurich, Switzerland, mmaes@ethz.ch

Andreas Flache

Department of Sociology/Interuniversity Center of Social Science Theory and Methodology, University of Groningen, 9712TG Groningen, The Netherlands, a.flache@rug.nl

Károly Takács

MTA TK “Lendület” Research Center for Educational and Network Studies (RECENS), H-1014 Budapest, Hungary, [karoly.takacs@uni-corvinus.hu](mailto:károly.takacs@uni-corvinus.hu)

Karen A. Jehn

Melbourne Business School, Carlton, Victoria, 3053 Australia, k.jehn@mbs.edu

Do strong demographic faultlines breed opinion polarization in work teams? We integrate two theories that have been used to explain faultline effects. The first, the approach of Lau and Murnighan [Lau DC, Murnighan JK (1998) Demographic diversity and faultlines: The compositional dynamics of organizational groups. *Acad. Management Rev.* 23(2):325–340], suggests that in teams with strong faultlines the mechanisms of homophilous selection of interaction partners and persuasive influence cause subgroup polarization, defined as the split of the team into subgroups holding opposing opinions. The second, from sociological and anthropological traditions, emphasizes that crisscrossing actors bridge faultlines because they share demographic attributes with several subgroups. Demographically crisscrossing actors help to prevent polarization in social groups. We argue that Lau and Murnighan’s theory implicitly factors in the effects of crisscrossing actors. However, we show that the authors overlooked crucial implications of their theory because they did not consider crisscrossing actors explicitly. Most importantly, we demonstrate that demographic crisscrossing implies that even teams with strong faultlines will overcome polarization in the long run, although they might suffer from it in the short term. We develop and analyze a formal computational model of the opinion and network dynamics in work teams to show the consistency of our reasoning with Lau and Murnighan’s theory. The model also revealed another counterintuitive effect: strong faultlines lead to structures of interaction that make teams with strong faultlines faster in arriving at a stable consensus than teams with weak faultlines.

Key words: group processes and performance; diversity in organizations; computer simulations; mathematical models; social networks

History: Published online in *Articles in Advance* September 10, 2012.

Introduction

Demographic diversity in the workplace is a major challenge for organizations and is becoming an increasingly important issue as the economy globalizes (for comprehensive reviews about theoretical and empirical research, see Bowers et al. 2000, Milliken and Martins 1996, Pelled 1996, Stewart 2006, van Knippenberg and Schippers 2007, Webber and Donahue 2001, Williams and O’Reilly 1998). For work teams, demographic diversity can be beneficial, because it broadens the social and human capital of the team. However, the benefits do not accrue automatically. Demographic dissimilarity between team members may, at the same time, cause conflicts and tensions and thus threaten performance. This leads Milliken and Martins (1996) to conclude in their review of the field that “diversity thus appears to be a double-edged sword” (p. 403).

In the search for conditions that explain why diversity sometimes increases team performance and reduces it at other times, Lau and Murnighan (1998, 2005) proposed that the effects of diversity may decisively depend on the way demographic attributes, like age and gender, are distributed among team members. Their main proposition is that diversity impairs team functioning when the distribution of demographic attributes generates a *strong faultline*: “Group faultlines increase in strength as more attributes are highly correlated, reducing the number and increasing the homogeneity of resulting subgroups. In contrast, faultlines are weakest when attributes are not aligned and multiple subgroups can form” (Lau and Murnighan 1998, p. 328). They argue that diversity (demographics not aligned) increases the potential of a team for creativity and good performance, but when the diversity is in a group with a strong faultline

(demographics aligned), this potential may not be realized. The team may split up into subgroups with polarized opinions that cause conflicts between team members (Bezrukova et al. 2009). An intriguing implication of faultline theory (Lau and Murnighan 1998) is that an ideal workgroup composition might exist such that large pools of social and human capital can be obtained, but the damaging effects of diversity on cohesion can be avoided.

We contribute to the faultline research by elaborating the explanation of the faultline proposition (Lau and Murnighan 1998) and thereby revealing crucial implications of faultline theory that have been overlooked so far. We start by reviewing two parallel theoretical lines that have been used to explain faultline effects and are based on fundamentally different arguments. The first, Lau and Murnighan's (1998) theory, highlights that the interplay of homophilous selection of interaction partners with social influence breeds subgroup polarization in work teams with strong faultlines. Subgroup polarization is our main dependent variable and is defined as the degree to which work team subgroups hold opposing opinions on work-related issues. The second theoretical line has been developed in classical sociological and anthropological studies and focuses on the integrating function of "crisscrossing actors" (Colson 1954, Evans-Pritchard 1939, Flap 1988, Galtung 1966, Lijphart 1977, Nieuwebeerta and Flap 2000, Ross 1920, Simmel 1908). Crisscrossing actors share at least one demographic attribute with another demographic subgroup than their own and thus function as a bridge over the subgroup split that was caused by the faultline.

Lau and Murnighan (1998) did not take into account how demographically crisscrossing actors can reduce subgroup polarization when the faultline is strong. Using their assumptions, we show that although their theory implicitly considers this, by not examining crisscrossing actors, the authors did not see some crucial implications of their theory. Foremost, we show that their theory predicts subgroup polarization only in the *short* term if the team contains demographically crisscrossing actors. We propose that in teams with a fixed composition, in the long run, demographically crisscrossing actors help to overcome subgroup polarization and group splits, even in teams with very strong faultlines. Moreover, we show that subgroup polarization in the short term depends on two further conditions that Lau and Murnighan implicitly assumed but which they did not examine theoretically. First, strong faultlines entail subgroup polarization in the short run only when employees exhibit sufficiently *strong homophily* when selecting interaction partners. That is, team members have a sufficiently strong preference for interacting with colleagues that are similar to them on certain attributes (McPherson et al. 2001). Second, we propose that demographic faultlines entail group splits in the short run only when there is sufficient initial congruency between work-related opinions and

demographic attributes in the work team (Homan et al. 2007, Phillips 2003, Phillips et al. 2004). In other words, we expect subgroup polarization only when initial opinion differences overlap sufficiently with the demographic faultline.

Faultline effects result from a complex interplay of the interactions between multiple team members. As Harrison et al. (2007) suggested, agent-based computational modeling is a powerful method that allows researchers to cope with theoretical complexity and to reveal counterintuitive implications of a theory. Frank and Fahrbach (1999) developed an agent-based model of complex and interrelated network and opinion dynamics in organizations, based on mechanisms that are very similar to those we assume. We follow their lead and use a formal modeling approach to show how Lau and Murnighan's reasoning, on the one hand, and the sociological and anthropological theories on crisscrossing actors, on the other hand, can be reconciled. We present and analyze a computational model that is based on Lau and Murnighan's core assumptions and test whether the propositions do consistently follow from their theory. Our results support the proposition that faultline effects occur only in the short run if the team comprises crisscrossing actors and that strong homophily and initial congruency are crucial conditions for the effect of faultlines on group polarization.

The analyses of our formal model also revealed a counterintuitive effect of faultline strength. Even though teams with a strong faultline tend to experience stronger subgroup polarization, our results suggest that they might be faster in arriving at a stable consensus once opinions have converged. It turned out that weak faultline teams overcome initial opinion differences quickly, but it also takes them longer to arrive at a stable consensus as team members tend to challenge the emerging consensus. This so-called "challenging process" (Lau and Murnighan 1998, p. 332) is shorter in teams with strong faultlines because these teams consist of subgroups that can quickly develop local consensus as a result of their internal coherence. Crisscrossing actors link these subgroups to each other such that opinion differences between groups gradually converge while their local coherence remains stable. Our analyses demonstrate that this process generates a stable outcome more quickly than the less structured consensus formation process in teams with weak faultlines, in which a larger number of mutual interactions between team members from different subgroups is possible.

Two Explanations of Faultline Effects

Lau and Murnighan's Explanation of Faultline Effects

Lau and Murnighan (1998) argued that all newly formed teams go through a "sensemaking process of understanding each other and their task" (p. 332) to coordinate

similar opinions about what their task is, how to fulfill it, and how to devise work. In this process, the interplay of two core mechanisms can cause problems in teams with a strong faultline. First, Lau and Murnighan assume *homophilous selection of interaction partners*. Team members tend to associate with colleagues who share relevant demographic attributes. This assumption is prominently supported by a large body of sociological research on homophily (Lazarsfeld and Merton 1954)—or the tendency of “birds of a feather flock together”—that has consistently been identified as a strong force in social interactions (McPherson et al. 2001). Studies in both educational (e.g., Kandel 1978, Moody 2001) and organizational settings (e.g., Bacharach et al. 2005, Ibarra 1992, McPherson and Smith-Lovin 1987, Ruef et al. 2003) have provided empirical confirmation of the homophily concept.

Homophilous selection of interaction partners implies that faultline strength affects who is interacting with whom in a team. To visualize this central point of Lau and Murnighan’s theory, we constructed six hypothetical teams with different faultline strengths. Each team consists of 20 individuals, a team size that is not too big to be unrealistic for a work team (Wegge et al. 2008) yet is also large enough to allow for a sufficiently fine-grained variation in the strength of demographic faultlines. Each team member is described by three dichotomous demographic attributes (symbolized as *black* versus *white*, *A* versus *B*, and *rectangle* versus *circle*). Diversity in all teams and on all demographic dimensions is kept at its maximum. That is, all three dichotomous variables have a distribution in which both values of the attribute are equally frequent (50:50). Teams differ, however, in the strength of their demographic faultlines. Faultline strength is denoted by the symbol f and is measured here by the pairwise Pearson correlation between all pairs of demographic attributes. Applying a method proposed by Flache and Mäs (2008a, b), we varied faultline strength between maximal (all pairwise Pearson correlations are 1) and minimal (all Pearson correlations are 0) faultline strength.

The network pictures in Figure 1 show the interaction structures of the six hypothetical work teams that follow from homophilous selection of interaction partners and the respective demographic faultline strengths. A team’s interaction structure is the social network that results from individual selection of interaction partners. The interaction structure shows to which degree team members form subgroups in the sense that there is frequent interaction within, and no interaction between, the subgroups. In the network pictures, team members are represented by nodes. To depict the effects of homophilous selection of interaction partners on the interaction structure, nodes have been arranged such that two individuals are placed nearer to each other the more demographic attributes they share (Kamada and

Kawai 1989, McFarland and Bender-deMoll 2007). Accordingly, if two nodes are placed near to each other, then the likelihood of interaction is high. In addition, we added a line between nodes that share at least one attribute to indicate that interaction between these individuals is likely. Whereas nodes without demographic similarity may also interact, and nodes with common demographic attributes may not always interact, demographic similarity ensures that there is always a positive probability of interaction between these nodes and that they are more likely to interact than two nodes without demographic overlap are.

The gray circles identify the biggest subgroups of maximally similar individuals. In the team with the strongest faultline (Team 1), the three demographic attributes correlate perfectly. Each pair of actors is either maximally similar or maximally *dissimilar*, and therefore either interacts frequently or never. On the team level, Figure 1 depicts two perfectly homogeneous but unconnected subgroups for Team 1. As faultline strength decreases, however, this separation between subgroups becomes weaker and completely disappears as faultline strength is minimal (Team 6). For instance, there are still two clearly distinct subgroups in Team 3 (medium faultline strength). However, the subgroups are smaller, and there are also team members that cannot be categorized into one of the subgroups. These actors share demographic attributes and therefore also interact with members of both subgroups.

In addition to homophily, Lau and Murnighan assume that during interaction, team members exert influence on each others’ opinions by *communicating arguments* (Isenberg 1986, Myers 1982, Myers and Lamm 1976, Vinokur and Burnstein 1978). They state, “Group members who support similar attitudinal positions will find that, as other members support that position using arguments different from their own, they each have more reason to become even more extreme than they were before” (Lau and Murnighan 1998, p. 332). Research on “polarization” (Myers 1982) has demonstrated how group members tend to become more extreme during group decision making. Faultline theory examines, however, not just polarization within one group but it focuses on what we denote as “subgroup polarization,” the degree to which a work team separates into subgroups *holding opposing opinions* (Lau and Murnighan 1998). Subgroup polarization during a team’s sensemaking process is problematic because it breeds emotional conflicts between the subgroups (Lau and Murnighan 1998), which in turn hamper good team performance (Jehn 1994, Jehn and Bendersky 2003, Pelled 1996).

The interplay of homophilous selection of interaction partners and influence with arguments can lead to subgroup polarization in groups with strong faultlines. As shown in Figure 1, homophily creates subgroups in teams with strong faultlines. Within subgroups, team

Figure 1 Interaction Structures of Six Hypothetical Teams with Different Faultline Strengths

members frequently communicate arguments, but argument communication *between* subgroups is rare. Lau and Murnighan argue that under these conditions, small initial opinion differences between the subgroups might be amplified during the sensemaking process. This is because subgroup members will mostly hear and share arguments that support their initial opinions (Stasser 1988), causing opinions in both subgroups to shift toward opposing ends of the opinion scale simultaneously. In other words, subgroup polarization increases. By contrast, in teams with weak faultlines, group members interact with colleagues who hold a variety of different opinions, such that no self-reinforcing dynamic toward emergent subgroup polarization can develop. In short, Lau and Murnighan's reasoning implies the following.

PROPOSITION 1. *The stronger the faultline in a work team is, the stronger subgroup polarization will be.*

Furthermore, the processes that Lau and Murnighan describe imply subgroup polarization only if two necessary conditions are met. First, the process of subgroup polarization crucially hinges on the assumption that *initial congruency* (Phillips 2003, Phillips et al. 2004) is sufficiently strong. That is, opinions and demographic attributes in a team need to be correlated initially, prior to interaction between the team members. If demographically similar group members do *not* share opinions more with each other than they do with demographically dissimilar others, then the exchange of arguments within demographic subgroups will not increase opinion differences between the groups. In this case members of one subgroup do not learn more new arguments pro or con toward the original opinion than the actors in the other subgroup. As a consequence, subgroup polarization will not occur. Thus, an initial correlation, or congruency, between demographic attributes and opinions appears

to be an essential condition for subgroup polarization even in work teams with a strong faultline. Although Lau and Murnighan's reasoning points to this condition, they have not made it explicit. Nevertheless, their theory implies the following.

PROPOSITION 1A. *Strong faultlines increase subgroup polarization only if initial congruency between demographic attributes and the opinion is sufficiently high.*

Second, subgroup polarization can only take place if homophily is sufficiently strong. We define the strength of homophily as the *degree* to which actors are more likely to interact with similar others than with dissimilar others (Lazarsfeld and Merton 1954, Wimmer and Lewis 2010). The strength of homophily in teams depends on team members' individual preferences for associating with similar colleagues, but it might also be determined by the institutional context of work teams. For instance, in teams with high task interdependence, workers are forced to collaborate with both similar and dissimilar colleagues to fulfill their tasks. Thus, in these teams similarity will only weakly influence the choice of interaction partners. As a consequence, team members interact frequently with members who hold different opinions and will thus be influenced by them. Such a context would make it unlikely that the teams' opinions polarize, even if faultlines are strong (Molleman 2005). Again, the ensuing proposition has been left implicit in previous theoretical research in faultline theory.

PROPOSITION 1B. *Strong faultlines increase subgroup polarization only if homophily is sufficiently strong.*

Crisscrossing as an Alternative Explanation of Faultline Effects

Almost a century ago, classical sociological and anthropological research on social order in stateless societies (Colson 1954, Evans-Pritchard 1939, Flap 1988, Galtung 1966, Lijphart 1977, Ross 1920, Simmel 1908) revealed that strong faultlines may cause a problem for social integration. For instance, Ross argued in his 1920 textbook:

Suppose at a given moment there is a certain strain along the line between Christians and Jews. If now, a strain appears along a quite different line, e.g., that between employers and workmen, the religious opposition will become less intense. For Jewish bosses and Jewish workmen will be estranged; likewise Christian bosses and Christian workmen. On the other hand, Jewish and Christian capitalists will recognize that they are "in the same boat," while Jewish workers and Christian workers will sympathize with one another as brother victims of exploitation.... Take the case of a tension between blacks and whites. Suppose now embitterment arises between labor and capital. If the lines of cleavage cross, each opposition will weaken the other. But if, as sometimes happens, all the employers are white men and all the employed are black men, then one antagonism helps the other and the rift in society is deeper than

ever.... A society, therefore, which is riven by a dozen oppositions along lines running in every direction, may actually be in less danger of being torn with violence or falling to pieces than one split along just one line. For each new cleavage contributes to narrow the cross clefts, so that one might say that society is *sewn together* by its inner conflicts.

(Ross 1920, pp. 164–165, italics in original)

Although both this classical sociological and anthropological literature and faultline theory agree on the prediction that strong faultlines breed conflicts, they base their prediction on different sets of assumptions. Whereas faultline theory argues that the interplay of homophily and social influence may result in subgroup polarization, the referred sociological and anthropological literature focuses on the integrating function of "crisscrossing" actors. Crisscrossing actors are individuals that share at least one demographic attribute with members of more than one demographic subgroup. Because of demographic similarity, they are attached to members of more subgroups and are thus able to conciliate in case of conflicts.

From this sociological and anthropological perspective, faultline effects follow because the more demographically crisscrossing actors there are in a group, the stronger are the integrating forces that prevent conflicts. The number of demographically crisscrossing actors in a group is, in turn, logically related to faultline strength. Figure 1 shows that the higher the number of those team members that are not part of one of the subgroups (i.e., crisscrossing actors), the weaker the faultline. Teams with the maximal faultline strength ($f = 1$, Team 1) consist of only two kinds of actors (*black*, *B*, *rectangles* and *white*, *A*, *circles*). There are no demographically crisscrossing actors in this team. The number of demographically crisscrossing actors increases as faultlines become weaker. Team 2, for instance, still consists of two large subgroups. However, there are also three crisscrossing actors present ($f = 0.8$). In short, the sociological and anthropological approach suggests the following.

PROPOSITION 2. *The more demographically crisscrossing actors there are in a group, the stronger are the integrating forces that prevent subgroup polarization.*

Integrating the Two Theories: Why Time Matters

The processes that the two explanations of the faultline hypothesis propose appear to be fundamentally different. On the one hand, Lau and Murnighan (1998) argue that in teams with strong faultlines, demographic subgroups form and develop increasingly different opinions that tear the team apart. The sociological and anthropological literature, on the other hand, points to those actors that connect the subgroups and prevent conflicts. We argue that it is of great importance for our understanding of faultline effects to analyze how exactly these two processes are related to each other. We have shown

that only groups with maximally strong faultlines have no demographically crisscrossing actors. If demographically crisscrossing actors can prevent group splits, does this then imply that their presence might neutralize the mechanism that Lau and Murnighan have described? Or would homophilous selection and argument communication undermine the integrating effects of crisscrossing actors if the faultline is sufficiently strong?

It turns out that the same assumptions from which Lau and Murnighan derive their faultline hypothesis can also be used to model the effects of crisscrossing actors. However, as we will show, when we explicitly integrate crisscrossing actors into Lau and Murnighan's reasoning, new consequences arise for the effect of faultline strength on the dynamics of subgroup polarization. With our integrating model, we can identify heretofore overlooked conditions under which the integrating effects of crisscrossing actors can be expected to prevail upon the dividing effect of a strong faultline.

In particular, we find that in teams with strong faultlines the processes that Lau and Murnighan describe breed polarization only in the early stage of the sensemaking process. Later, however, demographically crisscrossing actors will help overcome group splits. Homophilous selection implies that demographically crisscrossing actors interact with members of both subgroups because they have some demographic similarity with members of each group. Based on persuasive argument theory (Myers 1982), we can expect that they will get arguments from all sides and will also communicate them to all subgroups they interact with. In this way, demographically crisscrossing actors establish indirect communication between the subgroups who fail to interact directly. This may result in a gradual convergence of the subgroups' argument pools and also of their opinions, eventually reaching opinion consensus.

This view is inspired by insights from formal theories of social-influence dynamics in groups (Abelson 1964, French 1956, Harary 1959). These theories suggest that in the long run, even very few network links between otherwise unconnected subgroups suffice to decrease opinion differences between subgroups. Based on the assumption of Lau and Murnighan (1998) that demographic overlap implies interaction, demographically crisscrossing actors can be seen as the link that integrates all group members into the network of mutual social influences. This suggests that, in principle, one single crisscrossing actor might suffice to create enough indirect communication between two subgroups to bring their opinions together. Thus, even in a group with a strong faultline, a small number of demographically crisscrossing actors may ensure that no subgroup is entirely disconnected from outside influences. Accordingly, there should be *no* long-run effect of faultline strength on subgroup polarization, except for the extreme case of a maximally strong faultline that divides

the team into perfectly distinct subgroups. This absence of an effect of faultlines across almost the entire spectrum of possible teams is clearly contrary to what Lau and Murnighan (1998) suggest.

We propose that demographically crisscrossing actors help overcome group splits in the long run. However, we also argue that in the short run, the polarizing forces in teams with a strong faultline can be stronger than the integrative dynamic of indirect interaction through demographically crisscrossing actors. When homophily is strong, the members of the two subgroups will more likely interact with very similar subgroup members than with less similar crisscrossing actors. Accordingly, it is likely that within each subgroup a local opinion consensus develops. With high initial congruency, this consensus will likely be on an extreme opinion, and subgroups initially polarize (see Proposition 1). However, even though opinions have polarized, every member of the subgroup still has a positive probability of interacting with a demographically crisscrossing actor, at least from time to time. Whenever this happens, there is a chance that an argument from the out-group is adopted by in-group members. This argument can subsequently spread rapidly in the in-group because high demographic similarity within the in-group leads to frequent interaction between in-group members. In other words, we propose that the same mechanisms that, according to Lau and Murnighan, imply subgroup polarization in the short term also imply that subgroup splits are not stable in the long run if the group comprises demographically crisscrossing actors.

In sum, the integrated model suggests that a high initial congruency (Proposition 1A) and strong homophily (Proposition 1B) can give rise to subgroup polarization only in the short run. In the long run, all teams that comprise demographically crisscrossing actors will develop opinion consensus independent of the strength of homophily and initial congruency.

PROPOSITION 3. *In all teams where faultlines are not maximally strong, subgroup polarization occurs only in the short run. These teams will find consensus in the long run and will overcome subgroup polarization (assuming that the team composition remains fixed).*

We do not claim that long-run effects are always more important for work teams than short-run effects, especially because many work teams operate with real-time actions and deadlines and change their composition in the long run. We only emphasize that demographically crisscrossing actors cause long-run effects in teams with a fixed composition that are radically different from short-run effects.

The Model

The logical implications of the combination of homophilous selection, social influence, and faultline strength

result from a complex interplay of these mechanisms simultaneously operating in multiple actors responding to each others' behavior. The method of computational agent-based modeling (Adner et al. 2009, Harrison et al. 2007, Macy and Willer 2002) provides multiple examples of how in such a complex multiagent system simple theoretical assumptions may generate counterintuitive implications that would have been overlooked without model formalization (e.g., Frank and Fahrback 1999, Siggelkow and Rivkin 2006, Stasser 1988). Accordingly, we conducted a test of the logical consistency of our reasoning (Adner et al. 2009) with a computational model. Obviously, for the computational model, we needed to implement some simplifying assumptions. Whenever we felt some doubt that these assumptions are fully in line with the implicit assumptions of faultline theory, we conducted simulations also with alternative ways of operationalization. The central results of these additional analyses will be mentioned below and are discussed in more detail in the online appendix (available at <http://dx.doi.org/10.1287/orsc.1120.0767>).

The formal model is based on the two mechanisms that Lau and Murnighan's (1998) theory assumes: homophilous selection of interaction partners and social influence with arguments. Our analyses of model implications focus on the dynamics of team members' opinions and the degree to which opinion differences aggregate to subgroup polarization, the degree to which subgroups hold opposing opinions. However, the model also takes into account the arguments that form the basis of team members' opinions. In this model, each of the N team members is represented as an agent i characterized by D demographic attributes ($c_{i,d}$) and K opinions ($o_{i,k}$), where d refers to the d 'th demographic dimension and k to the k 'th opinion. The demographic attributes can take the value of either 1 or -1 ($c_{i,d} \in \{-1; 1\}$) and represent characteristics that are not open to social influence (e.g., gender, age) or that change at a much slower rate than opinions (e.g., skills, personal values). Therefore, we assume that the demographic attributes do not change in the progress of interaction. The opinions of the actors vary between -1 and $+1$ ($-1 \leq o_{i,k} \leq +1$) and are open to influence.

Following Lau and Murnighan's theory (1998, p. 332), we included that agents base their opinions on arguments $a_{k,l}$. For simplicity, we represent arguments as being either in favor of or against holding a pro opinion ($o_{i,k} > 0$) on the corresponding issue ($a_{k,l} \in \{-1; 1\}$). For each issue k , there exist P pro arguments ($a_{k,l} = 1$) and C con arguments ($a_{k,l} = -1$). Which arguments exist in a given work team setting is summarized in the arguments matrix. This matrix has K columns and $P + C$ rows. Cells with a row number smaller than $P + 1$ hold pro arguments; i.e., $a_{k,l} = +1$. The remaining rows hold con arguments; i.e., $a_{k,l} = -1$ (see Figure 2, matrix (a)).

Empirical research suggests that people have limited capacities to remember and process information (Cowan 2001, Miller 1956). Accordingly, we assume that agents base their opinion on S relevant arguments ($S \leq P + C$). Technically, an agent's opinion on issue k is the average value of the arguments ($a_{k,l}$) the agent considers as relevant (see Equation (1)). Thus, the more pro (con) arguments an agent's sample of arguments comprises, the higher (lower) the value of the agent's opinion will be. For simplicity, all S relevant arguments are equally weighted in the calculation of the opinion. This means that the relevant arguments do not differ in their persuasiveness. A technical implication of this assumption is that an agent's opinion can adopt only $S + 1$ different values:

$$o_{i,k} = \frac{1}{S} \sum_{l=1}^S a_{k,l}. \quad (1)$$

Following research on memory processes (Brown and Chater 2001), we assume that agents disregard arguments if they are not sufficiently recent. Thus, the more recent an argument is at a given point in time, the longer this argument will remain relevant for the formation of the agent's opinion. This is implemented for each agent in a relevance matrix. The relevance matrix has K columns and has $P + C$ rows. Each element indicates how recent the respective argument is for the agent. Elements of the relevance matrix with a row number smaller than $P + 1$ identify how recent pro arguments are, and the remaining elements determine how recent con arguments are. We denote how recent an argument is ($s_{l,i}$) with integer values between 0 and S ($s_{l,i} \in \{0, \dots, S\}$). A value of $s_{l,i} = 0$ indicates that the argument a_l is *not* sufficiently recent and therefore *not* relevant for actor i . Values above 0 indicate that this argument *is* sufficiently recent and therefore affects actor i 's opinion. The most recent argument has the value of $s_{l,i} = S$, the second-most recent argument has the value $S - 1$, and so on. Thus, if an agent considers three arguments ($S = 3$), then one has a recency of 1, one has a recency of 2, and one has a recency of 3. The recency rank of an argument does *not* affect the argument's persuasiveness—that is, the extent to which an argument shapes the current opinion (see Equation (1)). However, the recency determines *how long* an argument affects the agent's opinion in the influence process. The exact rules for updating argument recency will be elaborated further below (see the Argument Communication section).

Lau and Murnighan (1998) assume that a team's sensemaking process consists of a series of discussion meetings of the team's demographic subgroups (p. 332). This is a plausible scenario for small work teams, because subgroups consist of very few team members. For instance, a four-person team with maximal gender diversity consists of a male and a female dyad. However, in bigger teams (see, e.g., the teams in Figure 1), demographic subgroups consist of multiple members, and

meetings therefore require some sort of organization, an important aspect that Lau and Murnighan do not address. In contrast, empirical research showed that work-related issues are often discussed in informal dyadic interactions that provide a relevant channel of persuasion (Oh et al. 2004, Weenig 1999). We therefore model the sense-making process of a team as a sequence of informal dyadic interactions between team members in which thoughts about a work-related issue are communicated from one team member to the other. In small groups, this resembles Lau and Murnighan's assumption of subgroup meetings because demographic subgroups consist of very few team members. In bigger teams, modeling dyadic interaction does *not* exclude interactions like in a meeting, in which a team member communicates the same argument to a number of other members, but it breaks these interactions down in their most elementary unit, the transfer of an argument from one individual to another. In a similar vein, assuming dyadic interaction does not exclude mutual influence. In fact, it is very likely that two agents influence each other mutually because homophily implies that the selection of interaction partners is based on similarity, which by definition is a symmetric concept.

In sum, we model the sensemaking process of a team as a sequence of interaction events. Each interaction starts with the partner selection phase and is continued by the social influence phase. In the partner selection phase, two agents are matched for interaction, based on homophilous selection. Subsequently, an opinion of one of the interacting agents is updated, based on the argument communication mechanism.

Homophilous Selection

We implement the partner selection phase as follows. In each event, an agent i^* is randomly selected. All agents have at all events the same probability to be selected. Then an interaction partner j ($j \neq i^*$) is chosen. To incorporate homophily, the probability that actor j is chosen as interaction partner (p_j) depends on the similarity between i^* and j . In line with Lau and Murnighan's theory (see Figure 1) and empirical findings (Ibarra 1992, McPherson and Smith-Lovin 1987, McPherson et al. 2001, Ruef et al. 2003), we assume that similarity is based on the D demographic characteristics. In addition, empirical research following the similarity-attraction paradigm also suggests that opinion similarity increases the probability to interact (Byrne 1971). Therefore, Equation (2) includes that the similarity $sim_{i^*,j}$ between i^* and j depends on both demographic attributes and the opinions. We explored to which degree model implications depend on the consideration of opinion similarity in the selection process (see §6 of the online appendix). It turned out that all findings on the effects of demographic faultlines that we report in this paper can also be replicated with a

model version in which similarity is solely based on demographic characteristics and opinions are not taken into account. This suggests that considering opinion similarity in the selection process does not affect model implications in a critical way for the scenarios that we inspected.

Similarity $sim_{i^*,j}$ varies between 0 and 1. A similarity of 0 means that the two actors are maximally dissimilar, whereas a value of 1 indicates that both hold the same opinions and the same demographic attributes. We assume that all attributes are equally weighted in the calculation of similarity. Formally stated,

$$sim_{i^*,j} = \frac{1}{2 \cdot (D + K)} \left(\sum_{d=1}^D (2 - |c_{i,d} - c_{j,d}|) + \sum_{k=1}^K (2 - |o_{i,k} - o_{j,k}|) \right). \quad (2)$$

The probability that actor j is selected as an interaction partner of i (p_j) is derived from the *relative similarity* of i^* and j compared to the similarities of i^* to all other agents, except i^* herself. To vary the *strength of homophily*, we include, furthermore, a parameter h into the model ($h > 0$). The higher the value of h , the more the relative similarity of the focal agent i^* and agent j increases the probability that j will be chosen as an interaction partner:

$$p_j = \frac{(sim_{i^*,j})^h}{\sum_{j=1, j \neq i^*}^{N-1} (sim_{i^*,j})^h}. \quad (3)$$

In the computer program, the actual selection of i^* 's interaction partner was implemented in three steps. First, the interaction probability p_j is calculated for each of the $N - 1$ potential interaction partners ($j \neq i^*$) (Equation (3)). The interaction probabilities p_j sum up to 1. We partition the unit interval into subintervals such that to each potential interaction partner j a subinterval of length p_j is assigned. Second, a random number between 0 and 1 is drawn from a uniform distribution. Third, the agent j whose interval contains this randomly drawn number is selected for interaction. Note that the ordering of the agents in this sequence does not affect the outcome of the selection process. In this way, all agents with a nonzero interaction probability can be selected with a probability p_j . Thus, the more similar j is to i^* , the higher is the probability that they will interact. If two actors differ maximally with regard to their opinions and their demographic attributes, then the probability of interaction equals 0.

Argument Communication

After the interaction partners i^* and j^* have been selected for the respective event, agent i^* is influenced by j^* . We implement social influence through arguments in two steps. First, one of the arguments that j^* considers for the formation of her own opinion is adopted

by i^* . For this, one of the K opinions is selected randomly for update (k^*), with the same probability ($1/K$) for all opinions. Then, one of the S arguments that j^* considers relevant is picked (a_{k^*, l^*}) with equal probability ($1/S$) for all recent arguments. Arguments that are not relevant for j^* are not chosen. The chosen argument is adopted by i^* . Technically, the argument a_{k^*, l^*} in i^* 's relevance matrix adopts the value $S + 1$ ($s_{k^*, l^*, i^*} = S + 1$).

When an agent's relevance matrix has been updated repeatedly, it can happen that all existing arguments have been adopted at least once. However, it does not seem reasonable that after some time agents consider all arguments as relevant. In contrast, empirical research confirms that humans have very limited cognitive capacities and tend to disregard dated information (for a review, see Brown and Chater 2001; see also Sederberg et al. 2008). To take this into account, we implemented a second step of the influence process. The second step ensures that the number of arguments that are relevant for an agent remains constant at S during the whole sensemaking process. This implies that when an agent i^* has adopted an argument that has not been considered before, one of the arguments that is currently relevant for i^* will be dropped. We assume that agents drop the argument that has been adopted least recently. This reflects the idea that every time an agent hears an argument from an interaction partner, the cognitive importance of that argument is reinforced. The longer ago an argument has been heard from another agent without hearing it again, the less important the argument is considered to be, and sooner or later it will be seen as entirely unimportant (Brown and Chater 2001, Sederberg et al. 2008). Technically, we implement this in the model such that the relevance matrix of i^* is updated by subtracting one from all nonzero recency values with a higher recency than the transmitted argument prior to interaction. The argument that was communicated between i^* and j^* in the present event adopts a recency value of S ($s_{k^*, l^*, i^*} = S$) at the end of the iteration. All other relevant arguments decline in recency.

To illustrate the updating phase, Figure 2 contains two examples. Matrix (a) in Figure 2 is an example of an argument matrix with one column ($K = 1$) and four pro and four con arguments per issue ($P = C = 4$). Matrix (b) is the initial relevance matrix of agent i^* . Agent i^* bases her opinion on one pro and three con arguments ($S = 4$). According to Equation (1), this results in an opinion value of $o_{i,k} = -1/2$. Matrix (c) is the relevance matrix of i^* 's interaction partner j^* . Before the update, the first argument is not considered relevant by i^* (see the circle in matrix (b)), but it is considered relevant by j^* . Hence, it is possible that i^* adopts the first argument, resulting in the updated relevance matrix for i^* shown in (d). Here, the communicated argument is maximally recent (see the circle in matrix (d)). The recency of the remaining arguments has been reduced by 1. Note that this changed i^* 's opinion, which shifted from $-1/2$ to 0 because i^* adopted a pro and dropped a con argument. As a second example, assume now that not the first argument is selected for update, but argument number 5 instead. This argument has already been considered by i^* (see the square in matrix (b)). However, its recency has increased as a result of the interaction with j^* (see the square in matrix (e)). Note that this has no consequence on i^* 's opinion.

Lau and Murnighan (1998) did not specify which cognitive processes underlie social influence with arguments. Therefore, we tested alternative dropping rules to ensure that our findings do not critically depend on the assumptions we have added to the model. Most importantly, we implemented that the argument for dropping is selected at random instead of selecting the least recent argument. Computational experiments revealed that all qualitative results reported below are robust to this modification of the model (see §5 of the online appendix). This suggests that the results are not driven by the type of dropping rule we applied.

In addition, we experimented with the assumption that agents tend to drop those arguments that do not favor their current opinion (Heider 1967). We show in §5 of the online appendix that this is an additional mechanism

Figure 2 Example of the Updating Process with $K = 1$

that can generate subgroup polarization. However, the aim of this study is to test the exact implications of Lau and Murnighan's theory. Therefore, it is critical to avoid adding new mechanisms that interfere with Lau and Murnighan's theory. Hence, we excluded this assumption.

Interaction events are iterated until dynamics have reached equilibrium. In other words, after the opinion of agent i^* has been updated based on the interaction with j^* , the computer again picks an agent and selects an interaction partner that transmits an argument. This is repeated until no further changes in the distribution of arguments or opinions can occur in the dynamic system that the team constitutes. We deemed it necessary to run the simulation up to this point, because otherwise, misleading conclusions may be drawn about the determinants of consensus or polarization, as the short- or intermediate-run behavior of stochastic systems may be radically different from their long-run equilibrium behavior (Young 2001).

In our analyses, we study under which conditions the simulated work teams develop stable opinion consensus or split up into subgroups with opposing opinions. Nevertheless, to identify whether dynamics have reached a state of absolute stability, it is important to not only focus on the distribution of opinions but also take into account the distribution of arguments. For instance, if all team members hold the same opinion but base that opinion on different sets of arguments, then communication of arguments can still lead to opinion changes, and opinion consensus can break up again. This implication of argument-based influence was originally addressed by Lau and Murnighan (1998) in their informal theory, which holds that the sensemaking process of work teams may not be completed when the team has reached opinion consensus, because "one or more group members question the group's evolving, dominant script" (p. 332). This so-called "challenging process" is of particular importance during the sensemaking process of work teams, because "early group actions or decisions critically influence subsequent group processes" (Lau and Murnighan 1998, p. 332). We demonstrate below that the challenging process also plays a significant role in the simulated sensemaking process generated by the formal model of Lau and Murnighan's theory.

Our model has exactly two equilibria corresponding to perfect consensus or perfect subgroup polarization. Perfect consensus is reached when all agents hold the same opinions and base these opinions on the same arguments. In a work team that has reached this state, any further exchange of arguments will not lead to opinion changes because it will provide team members only with arguments that they already consider relevant. Perfect subgroup polarization is obtained if there are two subgroups, the members of each subgroup agree on all opinions and arguments with each other, and the pairwise similarity

($sim_{i,j}$) between agents of different subgroups is 0. That is, the members of the subgroups maximally differ with respect to all demographic attributes and all opinions. If all members of the subgroups base their opinions on the same arguments, then this outcome is stable. Any constellation that is not characterized by perfect consensus or perfect subgroup polarization is transient in the sense that it is logically possible that subsequent interaction leads to opinion changes of at least one team member. Therefore, the computer simulation process is continued until either perfect consensus or perfect subgroup polarization is reached.

Obviously, perfect subgroup polarization can only be stable in teams where faultline strength is maximal ($f = 1$), because in these teams there are no crisscrossing agents. Demographically crisscrossing agents share at least one demographic attribute with members of both demographic subgroups. Accordingly, if there is a demographically crisscrossing agent and the two subgroups still disagree, a positive probability remains that arguments of the one subgroup are adopted by the other and the disagreement will vanish.

Some of our analyses focus on the duration of the sensemaking process, i.e., the time that it takes before perfect consensus or perfect polarization has been reached. To be sure, we refrain from formulating statements about effects of the independent variables in our experiments on the *absolute* duration (e.g., in days or seconds) of the sensemaking process. We are not aware of any empirical evidence that would allow meaningfully assessing the duration of a simulated interaction event in real time. In addition, the duration of interaction events, and also their frequency per day or week, certainly depends to a considerable extent on the organizational and cultural setting, suggesting that auxiliary assumptions need to be included in order to arrive at testable conclusions about the exact duration of the sensemaking process in a given setting. Similar to Lau and Murnighan's theory, our model is not restricted to specific work team settings. Nevertheless, it seems reasonable to assume that the more interaction events occur before equilibrium, the longer such a process also would take in real time. This allows us to compare the *relative* length of the process in terms of number of simulation events under different conditions.

Simulation Experiments

The central outcome variable of faultline theory is the level of subgroup polarization in work teams. To quantify subgroup polarization, we use a measure called *polarization* (Flache and Mäs 2008a). It captures the degree to which the group can be separated into a small set of factions that are mutually antagonistic in the opinion space and have maximal internal agreement. To compute *polarization*, we use the variance of pairwise opinion agreement across all pairs of agents in the population, where agreement ranges between -1 (total

disagreement) and +1 (full agreement), measured as 1 minus the distance of opinions. This measure obviously adopts its lowest level of 0 for the case of perfect opinion consensus. The maximum level of subgroup polarization ($polarization = 1$) is obtained when the population is equally divided between the opposite ends of the opinion scale at -1 and $+1$, and all opinion dimensions are perfectly correlated.

To test whether the propositions follow consistently from the model, we conducted computational experiments, varying three model parameters: the strength of faultlines (f), the initial correlation between demographic attributes and opinions (w), and the strength of homophily (h).

Across all conditions of the simulation experiment, we assumed that the team members hold three salient demographic attributes ($D = 3$) and that diversity is maximal on all three dimensions ($= 50:50$ distribution). To allow under these specifications a sufficiently fine-grained variation of the central independent variable, faultline strength, we assumed a big but not unrealistic (Wegge et al. 2008) team size of 20 ($N = 20$). Furthermore, we assumed that the team focuses on one main issue ($K = 1$). All results could be replicated if two issues ($K = 2$) are taken into account, all other things being equal (see §4 of the online appendix).

Based on results from research on human's capacity to store information (Cowan 2001), we assigned the value 4 to S in all conditions. This assumes that the agents base their opinions on four arguments. We show in the online appendix that we could replicate our results also for a higher value of S . In general, we found that subgroup polarization was stronger when S was high. Finally, we assumed that there exist 10 pro ($P = 10$) and 10 con ($C = 10$) arguments. We selected values for P and C that are higher than S to create sufficient variation in the argument sets between pairs of agents with identical opinions. This decreases the chance that agents with similar opinions provide each other with arguments that they already consider relevant. Otherwise, interaction between agents with similar opinions does not lead to opinion changes, and the reinforcing effects of argument exchange that Lau and Murnighan describe cannot develop. $P = C = 10$ is high enough to avoid this and also reflects that the issue is not too complex to be unrealistic.

To vary faultline strength (f) in the experiments, we used exactly the same distributions of demographic attributes that we used in Figure 1 (Flache and Mäs 2008a, b). We varied the Pearson correlation between each pair of demographic attributes from 0 to 1 in steps of 0.2. Of course, there are many alternative distributions of the three variables that result in the same *bivariate* correlations. The distributions we used, however, have the property that they produce equal correlations between all pairs of demographic attributes and at the same time keep diversity maximal. We chose equal

correlations to resolve a conceptual unclarity in Lau and Murnighan's definition of faultline strength. Do we, for example, speak of a strong faultline if two variables x and y are perfectly correlated but completely unrelated to a third variable z ? Or, would we regard the faultline as stronger or weaker if x and y are correlated only with $r = 0.8$ but the correlation between x and z would rise to 0.6? These questions do not occur if all pairs or variables are equally correlated. Furthermore, considering unequal correlations between the demographic attributes does not affect the results. Also, with unequal correlations, it holds that the weaker the correlation between the demographic dimensions, the smaller the subgroups and the more crisscrossing actors there are in a team. Furthermore, as long as not all pairwise correlations are maximal ($f = 1$), crisscrossing actors are present.

To manipulate the level of initial congruency (w), we related the initial opinion to the first demographic attribute. The extent to which this affects the correlation of the opinion with the remaining demographic attributes depends on faultline strength (f). The stronger the faultline, the higher the correlation between the first demographic attribute and the other demographic attributes. Accordingly, the stronger the faultline, the more similar the correlations between the opinion and the first, second, etc., demographic attribute. Technically, we assigned S arguments to each agent. For each of the S arguments, we assigned one of the existing pro arguments with the probability w when the agent holds the value 1 at the first demographic attribute and one of the con arguments otherwise. Agents with the value -1 on the first demographic attribute received a pro argument with probability $1 - w$. For instance, if w is 0.5, then pro and con arguments always have the same probability to be assigned. On average, this results in a Pearson correlation between the first demographic attribute and the opinion of 0. However, as w increases, agents with the value 1 (-1) at the first demographic attribute more likely receive a pro (con) argument. This entails a higher Pearson correlation between the first demographic attribute and the opinion as w increases. Under $w = 1$, the opinion and the first demographic attribute perfectly align. More precisely, all agents that hold the value 1 at the first demographic attribute also hold opinion values of 1, and all agents who belong to the other demographic subgroup on the first dimension hold opinion values of -1 .

We varied w between 0.5 and 1 in steps of 0.1. We do not consider w -values below 0.5. Such values would lead to a negative correlation between the opinion and the demographic attributes. Because the actual values of the opinion and the demographic attributes have no substantial meaning, it makes no difference if opinions and demographic are positively or negatively correlated. To test the effects of the strength of homophily, we manipulated the parameter h (see Equation (3)), varying

it between 1 and 5 with steps of 1. A value of $h = 1$ expresses that agents have a weak preference to interact with similar team mates. The value of $h = 5$ corresponds to a very strong homophily.

All in all, we inspect $6 \cdot 6 \cdot 5 = 180$ conditions in our computational experiment. The five conditions in which faultline strength is maximal ($f = 1$) and the initial correlation between opinions and demographic attributes is maximal ($w = 1$) have been excluded because the similarity ($sim_{i,j}$) between agents is either 1 or 0 under this condition. In these cases, it is logically impossible that members of different subgroups will interact, and opinions cannot change. For reliability, we

conducted 500 independent replications per experimental condition.

Results

We present the results in three steps. In the first step, we present two ideal-typical simulation runs to illustrate the model dynamics. We then turn to the tests of the five propositions. Finally, we present additional analyses that revealed an unexpected effect of faultline strength.

Ideal-Typical Simulation Runs

Figure 3 demonstrates an ideal-typical simulation run with maximal faultline strength ($f = 1$). To trigger

Figure 3 Ideal-Typical Run with Maximal Faultline Strength ($f = 1$, $h = 5$, $w = 0.8$)

subgroup polarization, we imposed conditions that, according to the Propositions 1A and B, make polarization very likely. We assumed strong homophily ($h = 5$) and imposed a relatively strong correlation of initial opinions with demographic attributes ($w = 0.8$). The latter generated for this simulation run an initial Pearson correlation between the opinion and the three demographic attributes of 0.77. Figure 3 shows the development of *polarization* and the distribution of the opinion at different stages of the simulation run. The histograms show the respective opinion distribution. The network pictures describe the resulting interaction structure. In the network pictures, each agent is represented by a circle. The color (black or white) of a circle indicates to which of the two demographic subgroups the respective agent belongs. Each pair of agents that has a nonzero overall similarity ($sim_{i,j}$) is connected by a line, symbolizing the nonzero probability that these two agents interact. Because we focus here on the development of opinions in the team, the arrangement of the circles is only based on opinion similarity. Thus, circles are arranged in a way such that the nearer agents are placed to each other, the more similar their opinions are (Kamada and Kawai 1989, McFarland and Bender-deMoll 2007).

Initially (1st event in Figure 3), the opinion was almost uniformly distributed in this simulation run. Nevertheless, the corresponding network picture reveals that there are already initially systematic opinion differences between the demographic subgroups. The change of the histograms of the subsequent events shows that, over time, opinion differences between the subgroups increase. Consequently, the number of lines between the subgroups also decreases. Eventually (by the 3,400th event), the subgroups hold maximally opposing opinions. The exchange of arguments between subgroups stopped at this point because there is no overlap in demographic attributes or in opinions between agents from different subgroups. Opinion changes have now become impossible because agents only interact with team members who hold the same opinion and arguments. The online appendix contains further statistics of this ideal-typical run and an animation film (see the file “max_faultline.mov”) that illustrates the network dynamics.

Figure 3 depicts ideal-typical dynamics that ended in a stable group split. This shows that our model can generate the dynamics that Lau and Murnighan (1998) propose if the faultline is maximally strong. Proposition 3, however, expects that dynamics differ crucially when crisscrossing actors are present. Figure 4 shows an ideal-typical run that supports the proposition. In this run, faultlines were slightly weaker than in the condition of Figure 3. For comparison, we retain all further parameters of the first illustrative run with maximal faultline strength ($h = 5$, $w = 0.8$), but we slightly reduce the strength of the faultline to $f = 0.8$, which implies that

the team contains three crisscrossing agents (see the three squares in network pictures of Figure 4). Initially (see the 1st event), the opinion is again almost uniformly distributed and the demographic subgroups already hold somewhat different opinions. Again, we observe increasing subgroup polarization, just as Lau and Murnighan proposed. After the 705th event, the work team fell apart into two opposing subgroups with maximally different opinions. Within the two subgroups, the agents share the same opinions and also quickly coordinate on a common vector of arguments. However, the two subgroups are not completely unconnected: because of the three crisscrossing actors, there is still some exchange of arguments between the subgroups. The network picture of the 11,750th event demonstrates that one of the crisscrossing actors adopted an argument that changed his opinion. Subsequently, this argument spreads in the crisscrossing actors' subgroup, and the opinion differences between the subgroups decrease (see the 13,160th event). This convergence process continues until overall consensus is reached. In the online appendix, we provide an animation film (see the file “strong_faultline.mov”) of this ideal-typical simulation run.

Consistency Tests of the Propositions

Long-Term Effects. According to Lau and Murnighan (1998), higher faultline strength entails more subgroup polarization (Proposition 1). Proposition 3, however, claims that in teams with nonmaximal faultline strength, this effect can only be observed in the short term. Our experiments clearly confirmed Proposition 3. All simulated work teams with faultline strength below its theoretical maximum eventually ended in overall consensus. That is, all team members held the same opinion and based it on exactly the same arguments.

In teams with maximally strong faultlines, however, we found perfect subgroup polarization, but not in all simulation runs. Figure 5 shows how the initial congruency (w) and the strength of homophily (h) affected the frequency of runs that ended in perfect subgroup polarization. More precisely, the size of the bubbles in Figure 5 corresponds to the percentage of runs under the respective condition that ended in stable group splits. For instance, 49.6% of 500 runs with very strong homophily ($h = 5$) and very strong initial correlation between the opinion and the demographic attributes ($w = 0.9$) ended in a group split with two equally large subgroups and maximally opposing opinions. As suggested by Propositions 1A and 1B, the higher the values of h and w are, the more likely subgroup polarization occurs. Figure 5 thus confirms that our implementation of Lau and Murnighan's mechanisms can explain stable subgroup polarization in teams with maximally strong faultlines. However, even with maximal faultline strength, group splits remain unlikely if homophily is weak or the

Figure 4 Ideal-Typical Run with Three Crisscrossing Agents ($f = 0.8$, $h = 5$, $w = 0.8$)

opinions are not already initially strongly aligned with the demographic attributes. If homophily is weak, then agents exchange arguments across subgroup boundaries and thereby decrease subgroup polarization. If congruency is weak, then members of different demographic subgroups hold similar opinions that lead to communication between the groups and further convergence of opinions as well.

Short-Term Effects. Proposition 3 predicts increasing subgroup polarization in the short term even though teams reach consensus in the long term. To assess the short-term polarization in the simulated work teams, we measured the maximal value of *polarization* that teams exhibited during simulation runs and compared this value to the runs' initial value of polarization. The

difference between these two values indicates to which degree groups split up in the short run independent of whether the split occurred right at the beginning of the run or later. To test whether faultlines trigger short-term polarization (Proposition 3), Figure 6 shows bar graphs broken down by faultline strength (f). The gray part of each bar in Figure 6 depicts the average initial level of *polarization* in the teams. The black part of the bars shows the average increase in *polarization*. Both parts add up to the average of the maximal value of *polarization*. We excluded the conditions where $w = 1$, because here *polarization* is initialized at its logical maximum and cannot further increase. Figure 6 shows a stronger increase in the maximal value of *polarization* as faultlines become stronger (f) and thus supports

Figure 5 Percentage of Runs with Maximal Faultline Strength ($f = 1$) that Ended in Stable Splits

Proposition 3. At least in the short run, strong faultlines increase subgroup polarization.

According to Proposition 1A, the higher the initial congruency, the stronger the short-term effect of strong faultlines on polarization should be. To test that, we display in Figure 7 the effects of congruency (w) on maximal polarization broken down by faultline strength f . The figure shows that the initial polarization depends on w (see the gray parts of the bars). This is a technical consequence of congruency that occurs because opinions align closer with the 50:50 split on the values of +1 and -1 in the first demographic attribute, as w increases. There is, however, no such relationship of initial opinion polarization to faultline strength because the distribution of the first demographic attribute is the same for all levels of faultline strength. It turns out that the maximal value of polarization (see the complete bars) in all subgraphs increases with w . However, as the

Figure 6 Average Maximal Opinion Polarization Over Faultline Strength f (15,000 Runs per Bar)

size of the black areas shows, this is mainly the result of our manipulation of w . If faultlines are not strong ($f < 0.8$), the mean increase of polarization declines with the initial correlation between opinion and the first demographic attribute. We believe that this results from a ceiling effect. If faultlines are weak, then most pairs of agents have a relatively high similarity (sim_{ij}) because of shared demographic attributes. The potential of opinion polarization in these teams is thus very low. If w is high, these teams start out close to their potential maximum of polarization. As a consequence, polarization can only rise moderately above the initial level and will decline soon thereafter. If faultlines are strong ($f > 0.6$), however, the model produces the effect of w that Proposition 1A expected. The black parts of the bars in the subgraphs for faultline strengths of 0.8 and 1 show that a higher initial correlation of opinion and first demographic attribute entails more opinion polarization.

Proposition 1B suggests that subgroup polarization should increase with stronger homophily. Figure 8 confirms that the increase in polarization (see the black parts of the bars) is higher for stronger homophily (h). The comparison of different faultline levels also reveals that the effect of homophily strength increases in the strength of faultlines (f). If faultlines are weak, then even a very strong preference of the agents to interact with similar team members causes only a little increase in polarization in the short run. If faultlines are stronger, then strong homophily results in a larger increase in polarization.

Relative Time Until Convergence

The simulation experiments have confirmed that all teams that contained crisscrossing actors eventually arrived at overall consensus, even though many polarized in the short term. Teams where subgroup polarization increased in the short term but consensus was reached in the long run typically experienced subgroup polarization only at the beginning of the sensemaking process. As an illustration, in teams with maximal faultline strength ($f = 1$) and very strong homophily ($h = 5$) that eventually reached consensus, the highest degree of polarization that the respective team experienced was overcome, on average, after 9.5% of the overall duration of the respective simulation run.

The analyses of the length of the subsequent convergence process, however, led to an unexpected and counterintuitive result: the stronger the faultline in a team and the stronger homophily, the faster the teams arrive at overall consensus. Figure 9 shows a bubble graph expressing the average number of events it took until the runs ended in overall consensus, broken down by faultline strength (f) and homophily strength (h). The graph shows that the stronger the faultline was, the faster overall consensus was reached. It also shows that stronger homophily is associated with faster emergence of overall consensus.

Figure 7 Average Maximal Opinion Polarization Over Initial Congruency w , by Faultline Strength f (2,500 Runs per Bar)

To confirm the counterintuitive result about relative time needed until convergence, we conducted simulation experiments where teams started with perfect polarization ($w = 1$) and varied faultline strength (see

the online appendix). In the runs with weak faultlines ($f = 0$), the teams very quickly overcame the group split, but it took them very long to arrive at overall consensus. By contrast, it took the teams with a strong faultline

Figure 8 Average Maximal Opinion Polarization Over Homophily Strength h , by Faultline Strength f (2,500 Runs per Bar)

Figure 9 Average Number of Events Until the Teams Arrived at Equilibrium Based on the Same Arguments

($f = 0.8$) longer to overcome the initial group split. However, once the split was overcome, these teams quickly found a perfect consensus. Thus, teams with strong faultlines appear to find an overall consensus faster once a group split is overcome. This phase of the team process resembles what Lau and Murnighan (1998) refer to as the *challenging process*; in this phase, opinions have converged but “one or more group members question the group’s evolving, dominant script” (p. 332).

Why do teams with strong homophily and a strong faultline overcome the challenging phase faster? We explain this effect as a consequence of the interaction structure in teams with strong faultlines (see Figure 1). The same interaction structure that causes subgroup polarization in the short term accelerates the convergence process as soon as the opinion split has been overcome.

As we have shown in Figure 1, teams with strong faultlines consist of demographically homogeneous subgroups. If homophily is strong, agents mostly interact and exchange arguments with the members of their own subgroup. Within each subgroup, arguments that happen to be relevant for many subgroup members are communicated relatively frequently and are thus likely spread throughout the complete subgroup such that all subgroup members consider them relevant. The fewer subgroup members that consider a given argument relevant, the less likely this argument is to be communicated during interaction. As a consequence, arguments that are rather uncommon will not spread among subgroup members. Furthermore, those few agents that do consider an uncommon argument relevant are frequently exposed to alternative arguments. Therefore, their recency value

of the uncommon arguments will decrease and likely drop to 0. In this way, the overall number of arguments considered relevant by the members of each subgroup decreases over time, and eventually, members of each subgroup coordinate on a common set of arguments and find a local opinion consensus.

However, when there are crisscrossing agents, new arguments can enter a subgroup. It is possible that all subgroup members adopt a new argument and collectively drop one of the arguments used before. As a consequence, the subgroup members will coordinate on a new set of arguments and will find a new subgroup consensus on an opinion that is more similar to the other subgroup’s opinion than before the entry of the new argument. In addition, an argument that has been dropped by all subgroup members and that is not relevant for the members of the other subgroup will not reoccur in later interactions.

In teams with few crisscrossing agents (strong faultlines), new arguments enter the subgroups relatively seldomly. As a consequence, subgroups have enough time to coordinate on a common set of arguments before the next new argument enters. This leads to a *gradual* convergence of the opinions across subgroups. Instead, in teams with a weak faultline, there are more crisscrossing agents, and new arguments enter the discussion within subgroups more frequently. This can be so frequent that the members of the subgroups do not find consensus before a new argument enters. As a consequence, the number of arguments as well as opinion diversity within each subgroup remains high, and the gradual convergence of subgroups that we found in groups with strong faultlines does not develop. Furthermore, frequent argument exchange with crisscrossing actors leads to a fast spread of arguments across the entire team. Thus, if the members of a subgroup collectively drop an argument, this argument may still be used by other team members and might reenter the discussion in the subgroup over and over again. Overall, the convergence of opinions occurs faster in the interaction network of a team with a strong faultline than in the unstructured interaction pattern in a team with a weak faultline.

Summary and Implications

Lau and Murnighan (1998) argued that teams with strong demographic faultlines are likely to experience subgroup polarization. We challenged this prediction, arguing that Lau and Murnighan did not take into account the important role of demographically crisscrossing actors in the sensemaking process of teams. Demographically crisscrossing actors share some demographic attributes with multiple subgroups and can thus function as a bridge across the faultline. We showed that even teams with very strong faultlines comprise at least a few demographically crisscrossing actors and proposed that even teams with relatively strong faultlines

will eventually overcome polarization even though they might experience subgroup polarization in the short run. To underpin this claim, we developed a formal model based on the central behavioral assumptions of Lau and Murnighan's theory. We conducted computational experiments to test whether the new propositions follow consistently from the behavioral assumptions. Our analyses confirmed this.

In addition, we found that stronger faultlines imply opinion polarization only if demographic attributes are strongly correlated with the opinions of team members even before they influence each other and if team members select interaction partners based on strong homophily. Finally, contrary to intuition, our simulations revealed that teams with strong faultlines might be faster in arriving at opinion consensus. We found that teams with a strong faultline tend to experience subgroup polarization in the short run, but once opinions have converged, these teams relatively quickly find a stable consensus. Teams with a weaker faultline, on the other hand, experience less subgroup polarization at the beginning of the sensemaking process but go through a longer challenging process (Lau and Murnighan 1998), where team members question the evolving consensus.

Our analyses confirm that teams with strong faultlines experience more subgroup polarization than teams with weak faultlines. However, if faultlines are not maximally strong, as it is often the case in real-world teams, effects of faultline strength occur only for the short-term dynamics in a team. In the long run, group splits disappear sooner or later. This appears to be good news for managers. Nevertheless, we advise readers to interpret our results with caution. The main purpose of our analysis was to point to hidden implications of the mechanisms assumed by faultline theory. This does not preclude that other mechanisms not considered by the theory may lead to different consequences. Specifically, our formal model did not consider the possibility that social identities form around subgroups in the process of a group split (Rink and Jehn 2010). Members of the subgroups may then "act to legitimize the subgroups, and conflict between them may continue to be likely" (Lau and Murnighan 1998, p. 333). Strong subgroup identification may motivate team members to ignore arguments that support the opinion that is typical for the opposing subgroup or refuse interaction with crisscrossing actors (Abrams et al. 1990). Both would stabilize subgroup polarization. Identification might also promote the development of stereotypes about the demographic subgroups. Because crisscrossing actors fit into none of the stereotypes, they may be rejected by members of both demographic subgroups. If such negativity arises, then crisscrossing actors will not be able to conciliate.

Despite the possibility that identity formation may reduce the influence of crisscrossing actors, our results should not be discarded too readily. We have shown

that integrating effects of crisscrossing actors can in the long run only be precluded if these actors are perfectly excluded from the interaction networks within the subgroups. Even if subgroup identities form, it seems a rather extreme assumption that they can entirely prevent any subgroup influence via crisscrossing actors. It seems more plausible that the strength of subgroup identities affects how long it takes until the initial group splits can be overcome, but not the eventual outcome given that the team has enough time to converge to a perfect consensus. This also suggests that in the actual practice of work teams, crisscrossing actors may be important to overcome the negative effects of faultlines if management succeeds in creating conditions that support their integrating role. For example, an amicable and friendly environment in the work team may be important to reduce subgroup identifications and may therefore facilitate the exchange of arguments between the subgroups via crisscrossing actors (Jehn and Bezrukova 2010).

We recommend that future research on the logical implications of Lau and Murnighan's theory addresses two main areas. First, it would be interesting to investigate model predictions when additional mechanisms that may foster subgroup polarization, such as social identification processes, are included (see, e.g., §5 of the online appendix). This might help to identify conditions under which teams composed of crisscrossing actors do not overcome subgroup splits. Similarly, including additional assumptions about the frequency of interaction and the duration of argument exchange in a given work team setting might help to generate testable predictions about the *absolute* length of the sensemaking process.

Second, additional tests of robustness of our results are needed. Sections 4–6 of the online appendix summarize additional analyses that demonstrate robustness of results against assuming different parameter values and including alternative assumptions. One topic on which we focus in particular is the number of salient issues in the opinion space (K). We made the assumption that there was only one issue ($K = 1$) in the analyses of this paper, but teams may have to consider multiple opinion issues simultaneously. We expect, however, that a larger number of opinion issues will not fundamentally affect the dynamics of the model as long as Lau and Murnighan's assumption of strong initial congruency is retained. Strong initial congruency implies that all opinion dimensions are correlated with demographic characteristics. That is, for larger numbers of opinion dimensions, strong faultlines will give rise to initial opinion distributions in which there is a clear alignment of the demographic faultline with a group split in the multidimensional opinion space. From such a starting point, the mechanisms of homophily and argument exchange should amplify subgroup polarization independently of the number of opinion dimensions.

This is consistent with the robustness test that we conducted for $K = 2$. At the same time, the number of opinion dimensions should matter once the assumption of initial congruency is relaxed. To be more precise, if not all opinion dimensions are initially aligned with demographics, it becomes likely that even under a strong faultline, agents from different demographic subgroups come to interact, as a result of coincidental opinion similarities. The more opinion dimensions there are, the more there is some degree of opinion similarity between any two randomly chosen agents, and thus, the less we should see emergent polarization, all other things being equal. The online appendix contains more details on how different results are obtained when the assumption of strong congruency is relaxed.

The potential of coincidental opinion similarities to overcome demographic faultlines highlights a further direction for future research. Crisscrossing between subgroups may not only occur in demographic space, but it may also arise in opinion space. That is, even when a group falls apart in demographic subgroups, some actors may hold opinions, or develop them in the course of interaction, that give them sufficient overlap with members in each of the subgroups to communicate arguments in both directions. We have not explored this possibility systematically in our current analysis because we wanted to adhere as closely as possible to the original focus of Lau and Murnighan's theory, which is mainly on the role of demographic faultlines. However, the mechanisms of our model suggest that the presence of actors with crisscrossing opinions could have a similar integrating effect as demographic crisscrossing. The challenge is to identify conditions under which this could happen. Other than demographic features, opinions change in the course of interaction, and thus opinion crisscrossing is prone to be unstable. However, in a complex dynamic process, the increased likelihood of between-group interaction that opinion crisscrossing can bring about may be just the "push" that is needed to navigate the dynamic away from a state of stable polarization toward a state of consensus between the subgroups.

Our formal analysis of Lau and Murnighan's faultline theory illustrates that the method of agent-based modeling allows explicating complex theories and studying their implications over a wide range of conditions. However, agent-based simulation modeling can never explore the entire parameter space of a model, and therefore it does not allow proving robustness of results for all possible values of a given parameter, an endeavor that requires analytical approaches. However, considering the nonlinearity of central model mechanisms and the impact of randomness in our formalization, analytical modeling would presumably require making much more restrictive assumptions and developing more abstract models. It is a challenge for future research to explore whether analytical approximations are feasible without

sacrificing substantively essential features of the model proposed.

This paper revealed that short-term consequences of group dynamics might crucially differ from their effects in the long run. Other research has also proposed effects of time on consequences of demographic diversity in work groups. Most prominently, Harrison et al. (2002) argued that as team members get to know each other, the relevance of surface-level (demographic) characteristics will diminish, and members will base selection of interaction partners more on psychological similarity (personality, values, attitudes, beliefs). Like our reasoning, their argument suggests that the impact of demographic diversity and thus of demographic faultlines declines over time (see also Pelled et al. 1999). Surprisingly, our analyses have shown that this effect follows already from the elementary behavioral assumptions of faultline theory, without the need to necessarily include additional mechanisms such as the distinction between surface similarity and psychological similarity. This unexpected finding illustrates that already relatively simple models of social processes can be too complex to grasp their logical consequences by informal reasoning. Formal methods, therefore, are useful to study such complex systems and to reveal unexpected consequences of theories that may remain otherwise undiscovered.

Supplemental Material

Supplemental material to this paper is available at <http://dx.doi.org/10.1287/orsc.1120.0767>.

Acknowledgments

The authors thank Tobias Stark and André Grow, as well as the members of the Norms and Networks cluster at the Department of Sociology at the University of Groningen and the members of the Cooperative Relations and Social Networks Seminar at the Department of Sociology at Utrecht University for their constructive comments. The authors are grateful for many helpful comments by the two anonymous reviewers. M. Mäs, A. Flache, and K. Takács acknowledge financial support of this research by the Netherlands Organization for Scientific Research, NWO, under the Innovative Research Incentives Scheme [NWO/VIDI-Flache, Grant 452-04-351].

References

- Abelson RP (1964) Mathematical models of the distribution of attitudes under controversy. Frederiksen N, Gulliksen H, eds. *Contributions to Mathematical Psychology* (Holt, Rinehart Winston, New York), 142–160.
- Abrams D, Wetherell M, Cochrane S, Hogg MA, Turner JC (1990) Knowing what to think by knowing who you are—Self-categorization and the nature of norm formation, conformity and group polarization. *British J. Soc. Psych.* 29(2):97–119.
- Adner R, Polos L, Ryall M, Sorenson O (2009) Introduction to special topic forum. The case for formal theory. *Acad. Management Rev.* 34(2):201–208.
- Bacharach SB, Bamberger PA, Vashdi D (2005) Diversity and homophily at work: Supportive relations among white and African-American peers. *Acad. Management J.* 48(4):619–644.

- Bezrukova K, Jehn KA, Zanutto EL, Thatcher SMB (2009) Do workgroup faultlines help or hurt? A moderated model of faultlines, team identification, and group performance. *Organ. Sci.* 20(1):35–50.
- Bowers CA, Pharmer JA, Salas E (2000) When member homogeneity is needed in work teams. A meta-analysis. *Small Group Res.* 31(3):305–327.
- Brown GDA, Chater N (2001) The chronological organization of memory: Common psychological foundations for remembering and timing. Hoerl C, McCormack T, eds. *Time and Memory: Issues in Philosophy and Psychology* (Oxford University Press, Oxford, UK), 77–110.
- Byrne D (1971) *The Attraction Paradigm* (Academic Press, New York).
- Colson E (1954) Social control and vengeance in plateau Tonga society. *Africa: J. Internat. African Inst.* 23(3):199–212.
- Cowan N (2001) The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behav. Brain Sci.* 24(1):87–185.
- Evans-Pritchard EE (1939) Introduction. Peristiany JG, ed. *The Social Institutions of the Kipsigis* (Routledge, London), XIX–XXXIV.
- Flache A, Mäs M (2008a) How to get the timing right? A computational model of the effects of the timing of contacts on team cohesion in demographically diverse teams. *Comput. Math. Organ. Theory* 14(1):23–51.
- Flache A, Mäs M (2008b) Why do faultlines matter? A computational model of how strong demographic faultlines undermine team cohesion. *Simulation Model. Practice Theory* 16(2):175–191.
- Flap H (1988) *Conflict, Loyalty and Violence: The Effects of Social Networks on Behaviour* (Peter Lang, Frankfurt am Main, Germany).
- Frank KA, Fahrback K (1999) Organization culture as a complex system: Balance and information in models of influence and selection. *Organ. Sci.* 10(3):253–277.
- French JRP (1956) A formal theory of social power. *Psych. Rev.* 63(3):181–194.
- Galtung J (1966) Rank and social integration: A multidimensional approach. Berger J, Zelditch M, Anderson B, eds. *Sociological Theories in Progress* (Houghton Mifflin, Boston), 145–198.
- Harary F (1959) A criterion for unanimity in French's theory of social power. Cartwright D, ed. *Studies in Social Power* (Institute for Social Research, Ann Arbor, MI), 168–182.
- Harrison JR, Lin Z, Carroll GR, Carley KM (2007) Simulation modeling in organizational and management research. *Acad. Management Rev.* 32(4):1229–1245.
- Harrison DA, Price KH, Gavin JH, Florey AT (2002) Time, teams, and task performance: Changing effects of surface and deep-level diversity on group functioning. *Acad. Management J.* 45(5):1029–1045.
- Heider F (1967) Attitudes and cognitive organization. Fishbein M, ed. *Readings in Attitude Theory and Measurement* (John Wiley & Sons, New York), 39–41.
- Homan AC, van Knippenberg D, Van Kleef GA, De Dreu CKW (2007) Interacting dimensions of diversity: Cross-categorization and the functioning of diverse work groups. *Group Dynam. Theory Res. Practice* 11(2):79–94.
- Ibarra H (1992) Homophily and differential returns—Sex-differences in network structure and access in an advertising firm. *Admin. Sci. Quart.* 37(3):422–447.
- Isenberg DJ (1986) Group polarization: A critical review and meta-analysis. *J. Personality Soc. Psych.* 50(6):1141–1151.
- Jehn KA (1994) Enhancing effectiveness: An investigation of advantages and disadvantages of value-based intragroup conflict. *Internat. J. Conflict Management* 5(3):223–238.
- Jehn KA, Bendersky C (2003) Intragroup conflict in organizations: A contingency perspective on the conflict-outcome relationship. *Res. Organ. Behav.* 25:187–242.
- Jehn KA, Bezrukova K (2010) The faultline activation process and the effects of activated faultlines on coalition formation, conflict, and group outcomes. *Organ. Behav. Human Decision Processes* 112(1):24–42.
- Kamada T, Kawai S (1989) An algorithm for drawing general undirected graphs. *Inform. Processing Lett.* 31(1):7–15.
- Kandel DB (1978) Homophily, selection, and socialization in adolescent friendships. *Amer. J. Sociol.* 84(2):427–436.
- Lau DC, Murnighan JK (1998) Demographic diversity and faultlines: The compositional dynamics of organizational groups. *Acad. Management Rev.* 23(2):325–340.
- Lau DC, Murnighan JK (2005) Interactions within groups and subgroups: The effects of demographic faultlines. *Acad. Management J.* 48(4):645–659.
- Lazarsfeld PF, Merton RK (1954) Friendship and social process: A substantive and methodological analysis. Berger M, Abel T, Page CH, eds. *Freedom and Control in Modern Society* (Van Nostrand, New York), 18–66.
- Lijphart A (1977) *Democracy in Plural Societies: A Comparative Exploration* (Yale University Press, New Haven, CT).
- Macy MW, Willer R (2002) From factors to actors: Computational sociology and agent-based modeling. *Annual Rev. Sociol.* 28(1):143–166.
- McFarland D, Bender-deMoll S (2007) SoNIA (Social Network Image Animator). Accessed May 28, 2010, <http://www.stanford.edu/group/sonia>.
- McPherson M, Smith-Lovin L (1987) Homophily in voluntary organizations—Status distance and the composition of face-to-face groups. *Amer. Sociol. Rev.* 52(3):370–379.
- McPherson M, Smith-Lovin L, Cook JM (2001) Birds of a feather: Homophily in social networks. *Annual Rev. Sociol.* 27:415–444.
- Miller GA (1956) The magical number 7, plus or minus 2—Some limits on our capacity for processing information. *Psych. Rev.* 63(2): 81–97.
- Milliken FJ, Martins LL (1996) Searching for common threads: Understanding the multiple effects of diversity in organizational groups. *Acad. Management J.* 21(2):402–433.
- Molleman E (2005) Diversity in demographic characteristics, abilities and personality traits: Do faultlines affect team functioning? *Group Decision Negotiation* 14(3):173–193.
- Moody J (2001) Race, school integration, and friendship segregation in America. *Amer. J. Sociol.* 107(3):679–716.
- Myers DG (1982) Polarizing effects of social interaction. Brandstätter H, Davis JH, Stocker-Kreichgauer G, eds. *Group Decision Making* (Academic Press, London), 125–161.
- Myers DG, Lamm H (1976) The group polarization phenomenon. *Psych. Bull.* 83(4):602–627.
- Nieuwbeerta P, Flap H (2000) Crosscutting social circles and political choice—Effects of personal network composition on voting behavior in the Netherlands. *Soc. Networks* 22(4):313–335.
- Oh HS, Chung MH, Labianca G (2004) Group social capital and group effectiveness: The role of informal socializing ties. *Acad. Management J.* 47(6):860–875.

- Pelled LH (1996) Demographic diversity, conflict, and work group outcomes: An intervening process theory. *Organ. Sci.* 7(6):615–631.
- Pelled LH, Eisenhardt KM, Xin KR (1999) Exploring the black box: An analysis of work group diversity, conflict and performance. *Admin. Sci. Quart.* 44(1):1–28.
- Phillips KW (2003) The effects of categorically based expectations on minority influence: The importance of congruence. *Personality Soc. Psych. Bull.* 29(1):3–13.
- Phillips KW, Mannix EA, Neale MA, Gruenfeld DH (2004) Diverse groups and information sharing: The effects of congruent ties. *J. Experiment. Soc. Psych.* 40(4):497–510.
- Rink F, Jehn KA (2010) Divided we fall, or united we stand? How identity processes affect faultline perceptions and the functioning of diverse teams. Crisp RJ, ed. *The Psychology of Social and Cultural Diversity* (Wiley-Blackwell, Oxford, UK), 281–296.
- Ross EA (1920) *The Principles of Sociology* (Century Co., New York).
- Ruef M, Aldrich HE, Carter NM (2003) The structure of founding teams: Homophily, strong ties, and isolation among U.S. entrepreneurs. *Amer. Sociol. Rev.* 68(2):195–222.
- Sederberg PB, Howard MW, Kahana MJ (2008) A context-based theory of recency and contiguity in free recall. *Psych. Rev.* 115(4):893–912.
- Siggelkow N, Rivkin JW (2006) When exploration backfires: Unintended consequences of multilevel organizational search. *Acad. Management J.* 49(4):779–795.
- Simmel G (1908) *Soziologie: Untersuchungen über die Formen der Vergesellschaftung* (Duncker & Humbold, München, Germany).
- Stasser G (1988) Computer-simulation as a research tool—The discuss model of group decision-making. *J. Experiment. Soc. Psych.* 24(5):393–422.
- Stewart GL (2006) A meta-analytic review of relationships between team design features and team performance. *J. Management* 32(1):29–54.
- van Knippenberg D, Schippers MC (2007) Work group diversity. *Annual Rev. Psych.* 58(1):515–541.
- Vinokur A, Burnstein E (1978) Depolarization of attitudes in groups. *J. Personality Soc. Psych.* 36(8):872–885.
- Webber SS, Donahue LM (2001) Impact of highly and less job-related diversity on work group cohesion and performance: A meta-analysis. *J. Management* 27:141–162.
- Weenig MWH (1999) Communication networks in the diffusion of an innovation in an organization. *J. Appl. Soc. Psych.* 29(5):1072–1092.
- Wegge J, Roth C, Neubach B, Schmidt KH, Kanfer R (2008) Age and gender diversity as determinants of performance and health in a public organization: The role of task complexity and group size. *J. Appl. Psych.* 93(6):1301–1313.
- Williams KY, O'Reilly CA (1998) Demography and diversity in organizations: A review of 40 years of research. *Res. Organ. Behav.* 20:77–140.
- Wimmer A, Lewis K (2010) Beyond and below racial homophily: ERG models of a friendship network documented on Facebook. *Amer. J. Sociol.* 116(2):583–642.
- Young P (2001) *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions* (Princeton University Press, Princeton, NJ).

Michael Mäs is a postdoctoral researcher at the Chair of Sociology, in particular Modeling and Simulation, at ETH Zurich. He earned his Ph.D. from the Department of Sociology/ICS at the University of Groningen. His research focuses on collective action, intergroup conflict, integration in networks, and the emergence of social norms.

Andreas Flache is a professor of sociology at the Department of Sociology of the University of Groningen and the Interuniversity Center of Social Science Theory and Methodology (ICS) and is also a member of the board of directors. He received his Ph.D. from the University of Groningen. His research interests include applications of agent-based computational modeling and game theory to the study of cooperation and social integration, social network research, and research into social influence dynamics.

Károly Takács is the leader of MTA TK “Lendület” Research Center for Educational and Network Studies (RECENS) and an associate professor of sociology at the Corvinus University of Budapest, Institute of Sociology and Social Policy. He received his Ph.D. from the Department of Sociology/ICS at the University of Groningen. His main research interests are the theoretical and experimental analysis of the dynamics of social networks, in relation to problems of cooperation and conflict.

Karen A. Jehn is a professor of organization behavior at the Melbourne Business School. She earned her Ph.D. from Northwestern University. Her research focuses on intra-group conflict, group composition and performance, and lying in organizations. Her two most recent research interests are asymmetry of perceptions and member entitlement in workgroups.