How attitude certainty tempers the effects of faultlines in demographically diverse teams

André Grow · Andreas Flache

Abstract Lau and Murnighan's faultline theory suggests that strong demographic faultlines can undermine cohesion in work teams. A strong faultline splits a team into internally homogeneous but mutually dissimilar subgroups based on demographic characteristics. Social influence processes within these subgroups then lead to the polarization of team members' attitudes along the divisions imposed by the faultline. However, faultline theory hitherto neglects effects of attitude certainty. Research shows that the certainty with which individuals hold their attitudes affects social influence processes. We extend theoretical faultline research by integrating attitude certainty. For this, we incorporate the interplay of the dynamics of attitude certainty and social influence into a formal model of demographic faultline effects developed by Flache and Mäs. Computational experiments suggest a moderation effect. Demographic faultlines only affect team cohesion if attitude certainty is low. We discuss implications for future research.

A. Grow (✉) · A. Flache
Department of Sociology, ICS, Faculty of Behavioural and Social Sciences, University of Groningen, Grote Rozenstraat 31, 9712 TG Groningen, The Netherlands
e-mail: a.grow@rug.nl

A. Flache
e-mail: a.flache@rug.nl

1 Introduction

Recent decades saw an increase of both demographic diversity (Lau and Murnighan 2005; Tsui et al. 1992; van Knippenberg et al. 2004; van Knippenberg and Schippers 2007) and reliance on work teams (Jehn et al. 1999; Lau and Murnighan 1998; Milliken and Martins 1996) in organizations. The co-occurrence of these trends fostered researchers' interest in how diversity affects team processes and performance (Williams and Meân 2004; Williams and O'Reilly 1998). Two major perspectives emerged in the literature (van Knippenberg and Schippers 2007; Williams and O'Reilly 1998) that jointly suggest that diversity is a "double-edged sword" (Milliken and Martins 1996: 403). The *information/decision-making perspective* holds that diversity can benefit team performance, because it improves decision making through a broader basis of information and human capital in the team (van Knippenberg and Schippers 2007). By contrast, the *social categorization perspective* emphasizes that diversity may be detrimental to team performance, because it evokes social categorization and identification processes that threaten social cohesion and increase relationship conflict (van Knippenberg et al. 2004).

Given the downside of diversity, it appears questionable how teams can exploit its potential benefits (van Knippenberg et al. 2004). Lau and Murnighan's (1998) faultline theory opened up an intriguing new perspective on this problem. The theory suggests that demographic diversity does not necessarily reduce social cohesion. Rather, similar to Blau's (1977b) classical theory of crosscutting social circles, Lau and Murnighan pointed out that what matters is how multiple demographic characteristics (such as age and ethnicity) are correlated in a given team. Strongly correlated characteristics give rise to a *strong demographic faultline* that splits a team into few internally homogeneous but mutually dissimilar subgroups (e.g. a division between younger Hispanic women and older Caucasian men). They argued that under this condition team cohesion can come under pressure. Individuals preferentially interact with others whom they perceive as similar. Under a strong demographic faultline this can lead to segregated interaction in a team. *Social influence processes* then tend to reinforce preexisting differences in the *attitudes* of members of different demographic subgroups, entailing *subgroup polarization* along the demographic faultline. However, the outcome is different when demographic characteristics are uncorrelated, that is, when the faultline is *weak*. In such a situation, interaction between team members is not segregated and social influence processes are likely to promote *consensus* across different demographic subcategories in the team. Accordingly, Lau and Murnighan (1998) suggested that a team with weak demographic faultlines can benefit from its demographic diversity, whereas strong faultlines are likely to impede performance.

Empirical research in the wake of Lau and Murnighan's contribution has partially confirmed the expected negative effects of strong faultlines (Lau and Murnighan 2005; Molleman 2005; Thatcher et al. 2003). At the same time, the theory neglects hitherto an important insight from social psychological research on social influence that—as we argue in this paper—might yield new theoretical implications for the effects of faultlines on team cohesion. Faultline theory neglects that the extent to which individuals are open to social influence depends on their *attitude certainty*.

The less certain someone is about the correctness of her attitude, the more open she becomes for influence (Festinger 1950, 1954; McGarty et al. 1993; Smith et al. 2007). Conversely, the more certain a person is, the less susceptible she is to outside influence (Pomerantz et al. 1995). Social influence processes are in faultline theory the causal link between faultline strength and team cohesion. Hence, if we combine faultline theory with the notion that attitude certainty impedes influence, the implication should be that attitude certainty moderates the relation between faultlines and team cohesion. The higher the attitude certainty among team members, the less team cohesion should be affected by the strength of a demographic faultline.

The postulated moderation effect follows, however, less straightforwardly than our reasoning suggests up to this point. A second insight from social psychological research complicates the analysis. Empirical studies (e.g. Hensley and Duval 1976; Holtz 1997, 2003, 2004; Holtz and Miller 1985, 2001; Holtz and Nihiser 2008; Smith et al. 2007) indicate that attitude certainty is partially endogenous to social influence processes. That is, attitude certainty is itself subject to influence. Even individuals who hold their attitudes with high certainty can become less certain if nobody agrees with them. Similarly, individuals who hold their attitudes with low certainty can become more certain if others agree with them. This implies that any initial level of attitude certainty among team members might change in the course of social interaction. As a consequence, the degree to which attitude certainty impedes social influence is also subject to change.

The complex interdependency of attitude certainty, social influence, and faultline strength renders it difficult to gain solid intuitions about the aggregated outcomes of individual interactions in a team. Accordingly, in elaborating our theoretical argument we follow the lead of an increasing number of theorists who advocate the use of agent based computational modeling for understanding complex social dynamics with multiple interacting individuals (Bonabeau 2002; Harrison et al. 2007; Macy and Willer 2002; Smith and Conrey 2007).

Our contribution to faultline research is an integration of faultline theory with insights from research on the effects and dynamics of attitude certainty. We examine theoretically whether attitude certainty can moderate the effect of faultlines on team cohesion in terms of consensus and subgroup polarization. To this end, we draw on and extend an agent based computational model that Flache and Mäs (2008a, 2008b) proposed as a formalization of faultline theory.

In the remainder, we provide the theoretical basis of the assumptions of our model in Sect. 2. Next, the model is elaborated formally in Sect. 3. To explore model implications we conducted computational experiments. Results are reported in Sect. 4. Section 5 closes the paper with a discussion and outlook for future research. Note that in the social-psychological literature the terms *attitude certainty* and *uncertainty* are used interchangeably (e.g. Gross et al. 1995). To facilitate the model description, and following the convention in earlier models of effects of attitude certainty on influence processes (e.g. Deffuant et al. 2002), from here on we employ the terms *high* and *low* uncertainty which are equivalent to the terms low and high attitude certainty.

2 Theory

2.1 Demographic faultlines and team cohesion

Drawing on Blau (1977a) and O'Reilly et al. (1989), Lau and Murnighan (1998) assumed that individuals define their social identity to a large extent based on their demographic characteristics and have a preference to interact with those who are similar to them in these characteristics. Interaction in diverse teams is therefore likely to be structured by the distribution of demographic characteristics. However, Lau and Murnighan also stressed that individuals can hold multiple identities simultaneously (e.g. identities simultaneously based on gender, age, and ethnicity). Thus, the more several demographic characteristics coincide, the more likely social interactions will be constrained to demographic subgroups, because then different aspects of team members' identities induce the same patterns of interaction in the group.

The degree of alignment of several demographic characteristics is captured by the notion of demographic faultlines. An example helps clarifying this notion. Imagine two teams (cf. Lau and Murnighan 1998): Team A contains two Caucasian male members and two African-American female members; Team B contains one African-American and one Caucasian male member, as well as one African-American and one Caucasian female member. Both teams are alike in terms of diversity on each demographic characteristic considered separately. However, in Team A ethnicity and gender coincide and create a *strong* faultline that leads to clear-cut structures of non-overlapping subgroups based on both demographic characteristics. This facilitates identification processes and can lead interaction to occur mainly among the African-American male and among the Caucasian female team members. In Team B, ethnicity and gender create only *weak* faultlines, given that subgroup membership based on these characteristics is to some extent overlapping. Interaction is therefore less likely to be constrained to smaller subsets of team members.

Lau and Murnighan (1998: 332) suggested that the interaction structures that strong faultlines induce can lead subgroups to "find themselves polarizing and taking positions that become increasingly extreme". The reason is that similar individuals influence each other during their interaction by providing additional arguments that support each others' attitudes. Following a common social-psychological conception of attitudes, Lau and Murnighan assumed that attitudes toward an object (e.g. policies and procedures in an organization) can vary in degree from negative to positive (cf. Gross et al. 1995). Interaction within demographic subgroups reinforces pre-existing attitude similarity, whereas these attitudes become at the same time more extreme. That is, attitude variance within subgroups declines while attitudes shift at the same time toward either the negative or the positive side. In this process, reinforcement of preexisting within-group similarities increases preexisting attitude differences between subgroups, heightening the potential of intergroup conflict. However, if faultlines are weak interaction is not restrained to subgroups. Then, social influence more likely fosters overall consensus on salient issues.

As starting point of our formalization, we draw upon Flache and Mäs' (2008a, 2008b) recently proposed agent-based computational model of the effects of faultlines on team cohesion and extend their model to incorporate uncertainty in individuals' attitudes. Their model of social influence and selection builds on the four social

mechanisms *homophily*, *social influence*, *heterophobia*, and *rejection*. Homophily is the tendency of individuals to be attracted to, and to preferentially interact with, others who are like them (McPherson and Smith-Lovin 1987; McPherson et al. 2001; Milliken and Martins 1996). During interaction, similar individuals have the tendency to influence each others' attitudes (Festinger et al. 1950) and thereby become more similar and more attracted to each other (we refer to this also as *positive influence*). Heterophobia, on the other hand, is the tendency to avoid and dislike others whose dissimilarity exceeds a certain level (cf. Hensley and Duval 1976). If two individuals dislike each other, they have the tendency to reject each others' attitudes. That is, individuals modify their attitudes such that they become more dissimilar (we refer to this also as *negative influence*). This assumption is inspired by research building on the self-categorization paradigm in social identity theory (Brewer 1991; Hogg et al. 1990; Tajfel and Turner 1986) which holds that humans adjust their attitudes and behavior in a way to minimize the heterogeneity within their ingroups and to maximize differences to outgroups.

In this model, similarity perceptions between individuals are based on both demographic characteristics and attitudes. This is in line with Lau and Murnighan's (1998) assumption that also characteristics that are less obvious than surface demographic characteristics might determine individuals' similarity perceptions.

Flache and Mäs (2008a, 2008b) showed how their model generated two key results in line with Lau and Murnighan's theoretical expectations. First, teams with weak faultlines are likely to arrive at overall consensus. The reason is that if faultlines are weak, overall demographic similarity in the team is large enough to preclude strongly negative relations between most team members. Positive influence processes therefore prevail and eventually induce overall consensus. Second, teams with strong faultlines are likely to arrive at subgroup polarization. The reason is that if faultlines are strong, relations within demographic subgroups are likely to be positive (homophily), while relations between members of different demographic subgroups tend to be negative (heterophobia). As a consequence, positive influence occurs primarily within subgroups, whereas negative influence increasingly amplifies emerging attitude differences between subgroups. This leads to subgroup polarization.

We discuss next how integrating uncertainty in this model might alter these predictions.

2.2 Uncertainty and social influence

Uncertainty in attitudes refers to a situation in which individuals have low confidence in the correctness of their evaluation of an attitude object (cf. Festinger 1954; Gross et al. 1995; Holtz 2003; Smith et al. 2008). Uncertainty is central to a number of social psychological theories of influence processes (McGarty et al. 1993), such as the theory of social comparison processes (Festinger 1954) and self-categorization theory (Turner 1985). However, research varies in conceptualizing uncertainty as exogenous or endogenous to influence processes (cf. Gross et al. 1995).

2.2.1 Uncertainty as an exogenous factor

Theories that link uncertainty to individuals' openness to social influence typically conceptualize uncertainty as exogenous to influence processes. The more uncertain individuals are, the more open they become to influence by information that suggests a different evaluation of an attitude object. The reason is that they "have a fundamental need to feel certain about their world and their place within it [..., motivating] behavior that reduces subjective uncertainty" (Hogg and Mullin 1999: 253). Experiencing high uncertainty is an aversive state (Smith et al. 2007) and changing attitudes to make them more congruent with the provided information is a means to reduce uncertainty. However, individuals not only become more open to influence the more uncertain they are, they also become more resistant to it the more convinced they are that their attitudes are correct (McGarty et al. 1993; Turner 1991). This positive relation between uncertainty and influence found ample support in empirical research (e.g. Jetten et al. 2000; Pomerantz et al. 1995; Wilke et al. 1995).

The social-psychological notion of uncertainty is reflected in formal models of opinion dynamics that assume 'bounded confidence' (e.g. Deffuant et al. 2002; Hegselmann and Krause 2002). Following these models, we assume that the level of uncertainty translates into the size of a confidence interval that defines a range of 'acceptable' attitudes located within a certain distance around the focal individual's position on a continuous attitude scale. Only attitudes that fall within this range are perceived as 'plausible' or 'agreeable' and are therefore taken into account by the individual. The more uncertain the individual is, the wider this interval becomes. This conceptualization is mirrored in social judgment theory (Sherif and Hovland 1961) that assumes that individuals have a latitude of acceptance around their actual attitude that defines a range of positions that are considered acceptable. Also the size of this latitude has been assumed to vary with how confident an individual can be that her own attitude is correct (Insko et al. 1966).

Confining influence to attitudes within an acceptable distance imposes a new restriction on the dynamics that Flache and Mäs (2008a, 2008b) implemented in their model. While the overall similarity of two individuals determines whether they can influence each other positively or negatively, influence on a particular issue can actually only occur if there is sufficient attitude similarity on that issue. The less uncertain individuals are, the smaller is the range of attitudes that are deemed acceptable. As a consequence, even if two team members are in general very similar, they cannot influence each others' attitude on a particular issue if their attitude difference exceeds the threshold defined by their uncertainty. While this view implies that low levels of uncertainty impede both positive and negative social influence, it does not take into account that uncertainty itself can change during of the process of social influence. To this issue we turn next.

2.2.2 Uncertainty as an endogenous factor

Social-psychological research has shown that an individual's uncertainty is influenced by *informational/cognitive* factors, such as the knowledge about and experience with an attitude object, and by *social/consensual* factors, such as perceived

similarity of the own attitude with that of relevant others (Gross et al. 1995). Individuals typically experience high subjective uncertainty if they encounter new situations for which they have no relevant experience or insufficient information. Organization research has pointed to a range of settings that exhibit these characteristics. Examples are highly complex and dynamical organizational environments (Duncan 1972), organizations that switch from traditional hierarchical management to team based management (Bordia et al. 2004), or that introduce new and complex team tasks for which team members cannot rely on existing scripts and routines (Jehn 1995).

Perceived agreement or disagreement of the own attitude with that of relevant others has been shown to be an important social/consensual factor that influences uncertainty. Festinger (1950, 1954) argued that social/consensual factors become particularly important if informational/cognitive factors do not suffice to reduce uncertainty. Scholars of self-categorization theory go further and argue that social/consensual factors are the most important determinant of individuals' uncertainty, even in the presence of sufficient experience and knowledge (Hogg and Mullin 1999). They point out that the effects of agreement or disagreement with others crucially depend on the focal individual's perception of shared social identity.

On the one hand, uncertainty has been assumed to decrease if individuals find their cognitions in line with those of others with whom they identify, because these others are expected to "share key beliefs and values that are presumed to safeguard the integrity of their judgments" (Holtz 1997: 539). Conversely, if similar individuals disagree, uncertainty is likely to increase (McGarty et al. 1993; Smith et al. 2007), because disagreement indicates that one's attitudes are potentially incorrect. For both mechanisms researchers have found empirical support. A number of experiments conducted by Holtz and colleagues (Holtz 1997, 2003, 2004; Holtz and Miller 1985, 2001; Holtz and Nihiser 2008) suggest that already mere subjective perception of consensual support for one's attitudes by similar others is connected to decreased uncertainty. This effect is stronger the more similar attitudes of these others are assumed to be (Holtz and Miller 2001; Holtz and Nihiser 2008), and the more important the specific issue is to the individual (e.g. Holtz and Miller 1985, 2001). Furthermore, Smith et al. (2007) found that individuals report increased uncertainty if they are confronted with group norms that are ostensibly not in line with their point of view.

On the other hand, it has been assumed that the effect of attitudes of dissimilar others on uncertainty are antithetical to those of similar others. Agreement with dissimilar others should increase uncertainty because it challenges the social distinctiveness and thereby the correctness of the attitude of the focal individual (Hogg and Mullin 1999). However, if the attitudes of dissimilar others are clearly different, then the correctness of the attitude is verified in terms of social distinctiveness (Hogg and Mullin 1999) and uncertainty decreases (Holtz and Miller 2001). Indeed, Holtz and colleagues (Holtz 2003, 2004; Holtz and Miller 2001; Holtz and Nihiser 2008) found that the larger the perceived difference between the own attitude and that of dissimilar others, the less uncertain individuals are that their own attitudes are correct (see also Hensley and Duval 1976).

The research evidence leads us to incorporate into our model the assumption that a team member's current level of uncertainty increases to the extent that similar others disagree and dissimilar others agree with her. Conversely, we also assume that

uncertainty decreases if similar others agree and dissimilar others disagree with the team member. With these assumptions, effects of uncertainty on the relation between faultline strength and team cohesion do not only depend on the initial distribution of uncertainty in the team, but also on the dynamics that change uncertainty. For example, while initially low uncertainty may impede influence processes and thus prevent either consensus or polarization, team members' uncertainty may increase over time and give room to the dynamics that eventually entail either of these outcomes.

We translated these assumptions into a formal model, in order to obtain a solid intuition of how our extensions of faultline theory may change its implications. The model is fully elaborated in the next section.

3 The model

3.1 Team member characteristics and influence processes

Each of the N members of the virtual team is represented as an agent i, who is characterized by two types of attributes, namely demographic characteristics and attitudes. Demographic characteristics are 'fixed' and not subject to change, whereas attitudes are 'flexible' and subject to change. Agents have D demographic characteristics, where a_{id}^{fix} refers to the value of the dth demographic characteristic of agent i. Similarly, there are K issues in a team and $a_{ik,t}^{flex}$ is the attitude of i on the kth issue at time point t. Demographic characteristics can take the values -1 and 1 ($a_{id}^{fix} \in \{-1; 1\}$) and attitudes can vary continuously between -1 and 1 ($-1 \leq a_{ik,t}^{flex} \leq 1$). As indicated earlier, both demographic characteristics and attitudes determine the level of *dyadic similarity*, denoted $w_{ij,t}$, that agents perceive. We discuss below how this similarity is computed.

Each agent can hold each of her K attitudes with a different level of uncertainty, indicated by the *uncertainty level* $\mu_{ik,t}$. This parameter can vary continuously between 0 and 2 ($0 \leq \mu_{ik,t} \leq 2$) and creates an interval around the respective attitude that takes the form of $a_{ik,t}^{flex} \pm \mu_{ik,t}$. The interval indicates a range of values that i perceives as potentially correct. A larger interval indicates higher uncertainty. If $\mu_{ik,t} = 2$, the interval entails the entire attitude continuum, regardless of i's actual attitude.

The simulation algorithm proceeds as follows. Upon initialization of a simulation run, agents' demographic characteristics and attitudes are assigned and dyadic similarities are computed. Initial attitudes are randomly chosen from the entire interval of possible values, drawn from a uniform distribution. Initial demographic characteristics are used to impose an exogenously given strength of the demographic faultline, based on a procedure proposed by Flache and Mäs (2008a) that we explain below.

We express the length of a simulation run in iterations, where each iteration of the simulation consists of N steps. In each step one agent from the team is selected at random (with replacement) for either updating dyadic similarities ($w_{ij,t}$) *or* for

updating all attitudes ($a_{ik,t}^{flex}$) *and* all uncertainty levels ($\mu_{ik,t}$).[1] In the latter case, first, all attitudes are updated. After that, all uncertainty levels of the respective agent are updated.

The three updating processes are discussed in detail below. To express the temporal structure of updating dyadic similarities, attitudes, and uncertainty levels, we use in the following subscript t to refer to the state of a characteristic at the beginning of a particular updating process (i.e. at the beginning of updating dyadic similarities, updating of attitudes, or updating of uncertainty levels); subscript $t + 1$ indicates the state of a characteristic after the particular updating process is completed.

3.1.1 Updating dyadic similarities

Upon updating, dyadic similarity $w_{ij,t}$ is calculated as the average difference across both fixed and flexible attributes between agents i and j. Dyadic similarity can vary continuously between -1 and 1 ($-1 \leq w_{ij,t} \leq 1$). If $w_{ij,t}$ is equal to or larger than 0, i and j perceive each other as similar. If $w_{ij,t}$ is smaller than 0, they perceive each other as dissimilar. Dyadic similarity affects the strength and direction of the influence that another agent has on i's attitudes, and whether the agreement in attitudes between i and j increases or decreases i's uncertainty. We assume that weights are symmetric. Equation (1) formalizes the computation of dyadic similarity.

$$w_{ij,t+1} = 1 - \frac{\sum_{d=1}^{D} |a_{id}^{fix} - a_{jd}^{fix}| + \sum_{k=1}^{K} |a_{ik,t}^{flex} - a_{jk,t}^{flex}|}{(D + K)} \tag{1}$$

3.1.2 Updating attitudes and uncertainty levels

Upon updating, attitudes of the selected agent i are updated with regard to all K issues. For every issue updating is conducted in two steps. First, it is determined for all other agents j whether j can potentially affect i's attitude. The partial effect of a particular other agent j on i's attitude with regard to issue k, denoted $\gamma_{ijk,t}$, is mathematically defined in (2a) and (2b). The overall change of i's attitude is the sum of all *partial attitude effects* $\gamma_{ijk,t}$ from all other agents j. We inherit this assumption from a number of earlier models (Abelson 1964; Hegselmann and Krause 2002). The partial attitude effect of j on i's attitude is only different from 0 (i.e. $\gamma_{ijk,t} \neq 0$) if j's attitude $a_{jk,t}^{flex}$ is sufficiently similar to i's attitude, with boundaries given by $a_{ik,t}^{flex} \pm \mu_{ik,t}$. In this case, the magnitude of $\gamma_{ijk,t}$ indicates the strength of the effect that j's attitude has on that of i, which can vary continuously from 0 to 2 ($0 \leq \gamma_{ijk,t} \leq 2$).

Among similar agents, $\gamma_{ijk,t}$ is the difference between $a_{jk,t}^{flex}$ and $a_{ik,t}^{flex}$. In other words, if i considers j as similar, i's attitude is 'pulled' toward j's attitude by positive influence. For dissimilar agents, (2a) does not indicate the direction of influence, but only its magnitude. Equation (2b) implements, in combination with (3), the direction of influence. For dissimilar agents, this is away from the attitude of the source

---

[1] With probability 0.5 dyadic similarities are updated, otherwise attitudes and uncertainty levels are updated.

of influence. To obtain the magnitude of this negative influence in (2a), the attitude difference is subtracted from 2. As a consequence, the closer the attitude of two dissimilar agents, the stronger the attitude of agent j repulses i's attitude. This reflects that the more similar the attitude of an otherwise dissimilar agent is, the stronger is i's need for differentiation, as suggested by van Knippenberg and Wilke (1988).[2] As indicated in (2b), if two dissimilar agents hold the same attitude on a particular issue, the direction of the effect is chosen at random, given that there is no direction of repulsion implied.

$$
\gamma_{ijk,t} = \begin{cases} a_{jk,t}^{flex} - a_{ik,t}^{flex}, & \text{if } |a_{jk,t}^{flex} - a_{ik,t}^{flex}| \leq \mu_{ik,t} \wedge w_{ij,t} \geq 0 \\ (2 - |a_{jk,t}^{flex} - a_{ik,t}^{flex}|)\text{Sign}(x), & \text{if } |a_{jk,t}^{flex} - a_{ik,t}^{flex}| \leq \mu_{ik,t} \wedge w_{ij,t} < 0 \\ 0, & \text{if } |a_{jk,t}^{flex} - a_{ik,t}^{flex}| > \mu_{ik,t} \end{cases}
$$

(2a)

$$
\text{Sign}(x) = \begin{cases} 1, & \text{if } (a_{jk,t}^{flex} - a_{ik,t}^{flex}) > 0 \\ -1, & \text{if } (a_{jk,t}^{flex} - a_{ik,t}^{flex}) < 0 \\ random \in \{-1; 1\}, & \text{if } (a_{jk,t}^{flex} - a_{ik,t}^{flex}) = 0 \end{cases}
$$

(2b)

In a second step, actual attitude change is calculated. This is defined in (3). During the attitude updating process, i's new attitude, denoted $a_{ik,t+1}^{flex}$, is calculated as the sum of the original attitude and the weighted average of all partial attitude effects. The number of partial attitude effects is denoted $C_{k,t}$ (with $C_{k,t} \leq (N-1)$) and represents the number of cases in which i actually takes the attitude of another j into account (i.e. all cases for which $|a_{jk,t}^{flex} - a_{ik,t}^{flex}| \leq \mu_{ik,t}$). This number is multiplied by 2 in order to make attitude change more gradual. Each partial attitude effect is additionally weighted by $w_{ij,t}$. This implements two theoretical assumptions. First, the sign of $w_{ij,t}$ ensures that similar (dissimilar) others instigate a move of i's attitude toward (away from) j's attitude. Second, a specific difference in the attitudes of i and j has a stronger effect the more similar (dissimilar) they are in general. Equation (3) illustrates the basic principle of attitude change. Actual change is modeled as defined in (3a) and (3b) in order to avoid that attitudes take values smaller than -1 and larger

---

[2]This inversion is not implemented in the original model developed by Flache and Mäs (2008a, 2008b). However, this inversion is more in line with the notion that if dissimilar others hold very dissimilar (similar) attitudes, then the threat for the validity of the own attitude should be lower (higher), as put forward here. To check the comparability of the outcome of our model with that of the original model, we fixed the uncertainty attached to agents' attitudes to a value of 2. In this case, our model is equivalent to the model of Flache and Mäs, except for the rejection effect. The outcome that we obtained from our model was qualitatively identical to the outcome reported by Flache and Mäs.

than 1, and to smoothen attitude change at both ends of the attitude continuum.[3]

$$a_{ik,t+1}^{flex} = a_{ik,t}^{flex} + \frac{1}{2C_{k,t}} \sum_{i \neq j} w_{ij,t} \gamma_{ijk,t} \tag{3}$$

$$\Delta a_{ik,t}^{flex} = \frac{1}{2C_{k,t}} \sum_{i \neq j} w_{ij,t} \gamma_{ijk,t} \tag{3a}$$

$$a_{ik,t+1}^{flex} = \begin{cases} a_{ik,t}^{flex} + \Delta a_{ik,t}^{flex}(1 - a_{ik,t}^{flex}), & \text{if } \Delta a_{ik,t}^{flex} > 0 \\ a_{ik,t}^{flex} + \Delta a_{ik,t}^{flex}(1 + a_{ik,t}^{flex}), & \text{if } \Delta a_{ik,t}^{flex} \leq 0. \end{cases} \tag{3b}$$

Attitude updates are immediately followed by updates of the corresponding uncertainty levels, based on the location of i's new attitudes relative to those of all other agents. The new uncertainty level for a particular issue, denoted $\mu_{ik,t+1}$, is the sum of the original uncertainty and the weighted average of all effects that the attitudes of all other agents have on this uncertainty. This is mathematically defined in (4a) and (4b). The number of attitudes that impinge on i's uncertainty on the respective issue is denoted $A_{k,t}$ (with $A_{k,t} = (N-1)$) and is multiplied by 2 in order to make change in uncertainty more gradual. The size and direction of the effect that the attitude of any other agent can have on i's uncertainty depends on both whether i perceives j as similar or dissimilar and on the difference between $a_{jk,t}^{flex}$ and $a_{ik,t}^{flex}$ relatively to $\mu_{ik,t}$. For similar agents, uncertainty decreases if $a_{jk,t}^{flex}$ falls within $a_{ik,t}^{flex} \pm \mu_{ik,t}$. The smaller the attitude distance, the larger is the decrease of uncertainty. If $a_{jk,t}^{flex}$ does not fall into the interval $a_{ik,t}^{flex} \pm \mu_{ik,t}$, uncertainty increases with the distance between $a_{jk,t}^{flex}$ and the boundary of $a_{ik,t}^{flex} \pm \mu_{ik,t}$. For dissimilar agents, uncertainty increases to the extent that $a_{jk,t}^{flex}$ reaches into $a_{ik,t}^{flex} \pm \mu_{ik,t}$, but decreases to the extent that $a_{jk,t}^{flex}$ is distant from the boundary of $a_{ik,t}^{flex} \pm \mu_{ik,t}$. Each of these effects is weighted by $|w_{ij,t}|$ to implement the notion that j affects i's uncertainty more to the extent that i and j are either very similar or very dissimilar. To assure that uncertainty stays within valid boundaries, resulting values of $\mu_{ik,t+1}$ that fall outside of [0, 2] are reset to the nearest interval boundary.

These dynamics of uncertainty imply the possibility of changing asymmetries in influence within a dyad. For instance, if two agents are initially equally uncertain about their attitude about a specific issue (i.e. if $\mu_{ik} = \mu_{jk}$), differences in their own attitude relative to that of the surrounding agents can lead to differences in μ_{ik} and

---

[3]Equation (3) differs from (1) in Flache and Mäs (2008b) in two aspects. First, Flache and Mäs model the effect that the attitude of another agent j can have on i's attitude as $w_{ij,t}(a_{jk,t}^{flex} - a_{ik,t}^{flex})$ (note that we added the subscript t at appropriate places). Hence, in the original model, i's attitude is affected by the attitudes of *all* other agents, weighted by their similarity to i. We model this effect by $w_{ij,t}\gamma_{ijk,t}$. This means that in our model i's attitude is not necessarily affected by the attitudes of all other agents, but *only* by those attitudes that fall within $a_{ik,t}^{flex} \pm \mu_{ik,t}$. Second, in Flache and Mäs $C_{k,t}$ is always equal to the number of all j's in the team who are linked to i in the access network. Here $C_{k,t}$ is equal to the number of j's whose attitudes fall within $a_{ik,t}^{flex} \pm \mu_{ik,t}$.

μ_{jk} (e.g. j might become less uncertain than i), possibly resulting in an asymmetric influence relation between i and j.

$$\mu_{ik,t+1} = \mu_{ik,t} + \frac{1}{2A_{k,t}} \sum_{i \neq j} |w_{ij,t}| \Delta \mu_{ik,t} \tag{4a}$$

$$\Delta \mu_{ik,t} = \begin{cases} (|a_{jk,t}^{flex} - a_{ik,t}^{flex}|) - \mu_{ik,t}, & \text{if } w_{ij,t} \geq 0 \\ \mu_{ik,t} - (|a_{jk,t}^{flex} - a_{ik,t}^{flex}|), & \text{if } w_{ij,t} < 0 \end{cases} \tag{4b}$$
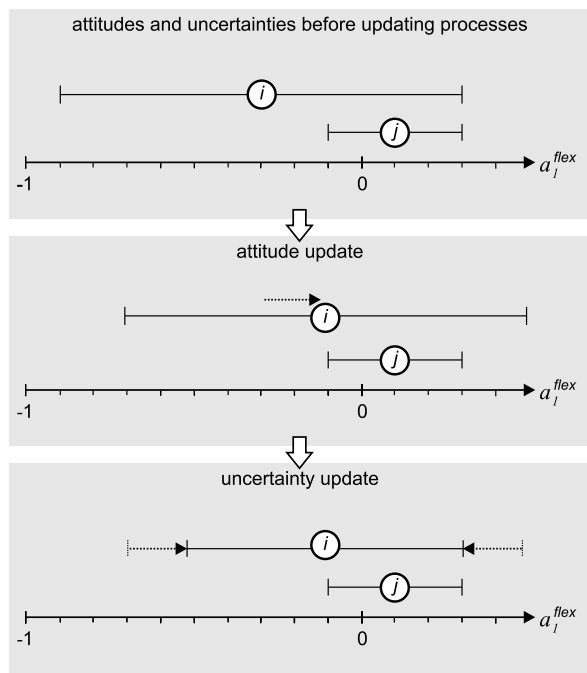
Figures 1 and 2 illustrate the mechanisms defined in this section. Both figures represent a simplified situation with a team that consists of two agents, i and j. In Fig. 1, both agents are maximally similar on all fixed and flexible attributes, except for a_1^{flex}; in Fig. 2, both agents are maximally dissimilar (similarity is indicated by shading of the circles). The attitude values of both agents regarding a_1^{flex} are indicated by the position of the circles along the x-axis. The uncertainty level attached to both agents' attitudes is represented by error-bars.

In part (a) of Fig. 1, the attitudes of the agents before the updating processes are $a_{i1}^{flex} = -0.3$ and $a_{j1}^{flex} = 0.1$. The uncertainties attached to these attitudes are $\mu_{i1} = 0.6$ and $\mu_{j1} = 0.2$. Assume now that agent i is selected for updating attitudes and uncertainty levels. Agent i is positively influenced by j, because $a_{j1,t}^{flex}$ falls within $a_{i1,t}^{flex} \pm \mu_{i1,t}$, and their mutual similarity weight $w_{ij,t}$ is positive. As a consequence i's attitude shifts towards j's attitude. That is, $a_{i1,t+1}^{flex}$ is closer to $a_{j1,t}^{flex}$ than $a_{i1,t}^{flex}$ (note that from t to $t + 1$ agent j's attitude does not change). After updating i's attitudes the uncertainty levels are updated. Agent i's uncertainty in $a_{i1,t}^{flex}$ decreases (i.e. the bars become narrower), because $a_{j1,t}^{flex}$ falls within $a_{i1,t}^{flex} \pm \mu_{i1,t}$ and $w_{ij,t} > 0$ (note that from t to $t + 1$ agent j's uncertainty level does not change). Part (b) of Fig. 1 illustrates a case in which the updating of i's attitudes already took place; we are thus only interested in changes in i's uncertainty level. In this case i is initially less uncertain than in the case shown in part (a). As the increase in the size of μ_{i1} indicates, i becomes more uncertain in a_{i1}^{flex}, because $a_{j1,t}^{flex}$ does not fall within $a_{i1,t}^{flex} \pm \mu_{i1,t}$.
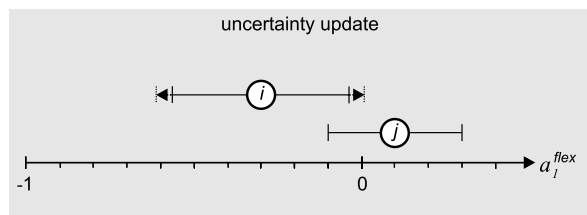
Part (a) of Fig. 2 shows the same attitude configuration that was shown in part (a) of Fig. 1. This time it is assumed that i and j are dissimilar, with a negative mutual similarity weight $w_{ij,t}$. In part (a) of Fig. 2, i is negatively influenced by j. As a consequence, the new attitude $a_{i1,t+1}^{flex}$ is more distant from $a_{j1,t}^{flex}$ than $a_{i1,t}^{flex}$. Subsequently, agent i's uncertainty in a_{i1}^{flex} decreases, because $a_{j1,t}^{flex}$ does not fall within $a_{i1,t}^{flex} \pm \mu_{i1,t}$ and $w_{ij,t} < 0$. Again, part (b) of Fig. 2 illustrates a case in which the updating of i's attitudes already took place. As the increase in the size of μ_{i1} indicates, i becomes more uncertain in a_{i1}^{flex}, because $a_{j1,t}^{flex}$ falls within $a_{i1,t}^{flex} \pm \mu_{i1,t}$.

Fig. 1 Changes in attitude and uncertainty of agent i in a_1^{flex} in the presence of agent j ($N = 2$); the agents are maximally similar on all other attributes (i.e. $w_{ij,t} = 0.94$). Error-bars attached to the attitude position of each agent illustrate the uncertainty level attached to the respective attitude


3.2 Conditions for the computational experiments

We vary two key conditions between teams in our computational experiments, strength of demographic faultlines and the level of initial uncertainty among team members. We assume teams with 20 members ($N = 20$), each with 3 demographic characteristics ($D = 3$) and 4 attitudes ($K = 4$). Using 3 demographic characteristics enables us to vary faultline strength sufficiently fine-grained for the purpose of this study. At the same time, using 4 attitudes ensures that agents' perceptions of mutual similarity/dissimilarity are not fully determined by demographic attributes. Similar to Flache and Mäs (2008a, 2008b) we implement 6 levels of faultline strength as shown in Table 1. In this table, each row represents one agent and each column corresponds

Fig. 2 Changes in attitude and uncertainty of agent i in a_1^{flex} in the presence of agent j ($N = 2$); the agents are maximally dissimilar on all other attributes (i.e. $w_{ij,t} = -0.77$). Error-bars attached to the attitude position of each agent illustrate the uncertainty level attached to the respective attitude

a – agreement before updating attitudes

attitudes and uncertainties before updating processes

attitude update

uncertainty update
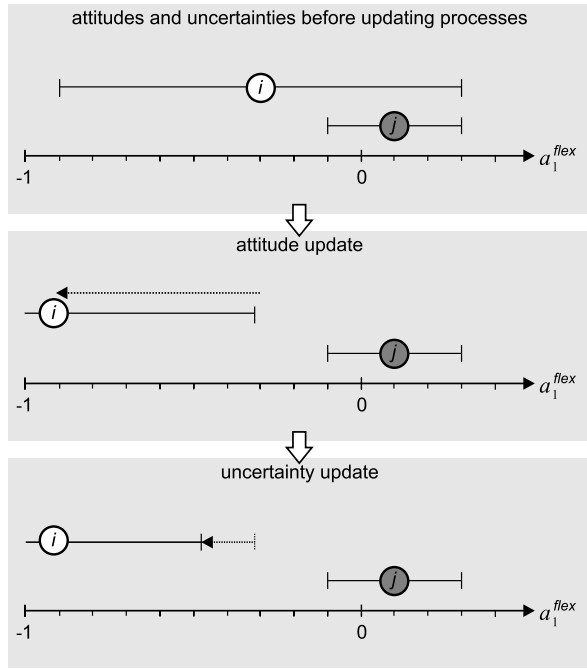
b – disagreement before updating uncertainties

uncertainty update
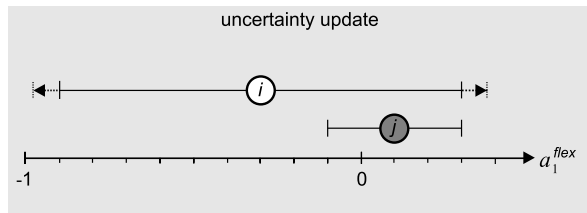
to the value on a specific demographic characteristic for the respective agent. The frequency of each of the two possible values of each demographic characteristic—which are differentiated by shading to facilitate visual identification—is the same in each column. In this way, diversity in each characteristic is constant within and across faultline conditions. This enables us to disentangle effects of faultline strength from effects of different levels of diversity. Faultline strength, denoted r, is expressed as the correlation between the three demographic characteristics. If $r = 1$, all three demographic characteristics coincide across all agents. In this case, agents can be divided into two equally sized subgroups whose members are internally maximally similar but mutually maximally dissimilar in terms of demographic characteristics. With decreasing r, subgroup boundaries become more diffuse.

Table 1 Implementation of different levels of faultline strength

i	r = 0.0			r = 0.2			r = 0.4			r = 0.6			r = 0.8			r = 1.0		
	d1	d2	d3	d1	d2	d3	d1	d2	d3	d1	d2	d3	d1	d2	d3	d1	d2	d3
1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
2	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
3	-1	-1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
4	-1	-1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
5	-1	-1	1	-1	-1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1
6	-1	1	-1	-1	-1	1	-1	-1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1
7	-1	1	-1	-1	1	-1	-1	-1	1	-1	-1	-1	-1	-1	-1	-1	-1	-1
8	-1	1	-1	-1	1	-1	-1	1	-1	-1	-1	1	-1	-1	-1	-1	-1	-1
9	-1	1	1	-1	1	1	-1	1	-1	-1	1	-1	-1	-1	1	-1	-1	-1
10	-1	1	1	-1	1	1	-1	1	1	-1	1	1	-1	1	-1	-1	-1	-1
11	1	-1	-1	1	-1	-1	1	-1	-1	1	-1	-1	1	-1	-1	1	1	1
12	1	-1	-1	1	-1	-1	1	-1	-1	1	-1	1	1	1	1	1	1	1
13	1	-1	-1	1	-1	1	1	-1	1	1	1	-1	1	1	1	1	1	1
14	1	-1	1	1	-1	1	1	1	-1	1	1	1	1	1	1	1	1	1
15	1	-1	1	1	1	-1	1	1	1	1	1	1	1	1	1	1	1	1
16	1	1	-1	1	1	-1	1	1	1	1	1	1	1	1	1	1	1	1
17	1	1	-1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
18	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
19	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
20	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
$\sum$	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

Table 2 Different levels of initial uncertainty (μ_{in}) and corresponding uncertainty level (μ_i)

μ_{in}	μ_i
0.0	0.0
0.2	0.4
0.4	0.8
0.6	1.2
0.8	1.6
1.0	2.0

Teams differ, furthermore, in the level of *initial uncertainty*, denoted μ_{in}, that defines at the outset the same size of the uncertainty level for all agents and all attitudes. For ease of interpretation, we let values of the parameter initial uncertainty vary continuously from 0 to 1 ($0 \leq \mu_{in} \leq 1$), where 1 indicates that all agents have initially the largest possible uncertainty level ($\mu_i = 2$) and 0 indicates the lowest possible uncertainty level ($\mu_i = 0$) for all agents. We investigate the effects of 6 different levels of initial uncertainty. Table 2 shows how each of these levels relates to numerical realizations of the uncertainty level. To disentangle general effects of the level of initial uncertainty in a team from effects that might arise from differences in uncertainty among team members, all team members are initially equally uncertain in all their attitudes. However, these uncertainties can independently change in the course of the simulation process.

3.3 Aggregate outcome measures

Our main interest is in studying the effects of the strength of demographic faultlines and initial uncertainty on the degree of consensus or polarization in the distribution of attitudes that results in a team. We use two measures to assess the level of polarization of demographic subgroups, *polarization* and *co-variation of fixed and flexible attributes*. In addition, we use the measure *attitude variance* to assess the degree of consensus in a team. To indicate the degree to which, overall, agents are open for social influence, we also measure *average uncertainty* across all team members. The measures of the co-variation of fixed and flexible attributes and the measure of the attitude variance are adopted from Flache and Mäs (2008a, 2008b).

The measure of *polarization*, denoted *pol*, captures the extent to which the team falls apart into subgroups such that agents' attitudes are maximally similar within a subgroup but maximally dissimilar between the subgroups. A team is maximally polarized if it is split into two equally large such subgroups. The measure is obtained in two steps. In the first step, we identify the subgroup of agents that most clearly forms a cohesive subgroup in terms of its internal agreement and disagreement with those agents who do not belong to this subgroup. In the second step, we obtain a measure that captures the degree to which internal attitude agreement within this subgroup exceeds agreement between members and non-members of the subgroup.[4]

To find the most cohesive subgroup we adapt the concept of *LS* sets developed in graph theory. The "definition of an *LS* set compares the frequency of ties within and between subsets [in a graph ...]. Lines within the *LS* set [...] should be more numerous than lines from a subset of nodes in an *LS* set to non-*LS* set members" (Wasserman and Faust 1994: 269). The graph on basis of which we identify cohesive subsets is given by the pairwise attitude similarities, denoted $s_{a^{flex}ij}$. Technically, $s_{a^{flex}ij}$ is obtained as two minus the absolute value of the average difference in agents' attitudes, as given by (5).

$$s_{a^{flex}ij} = 2 - \frac{\sum_{k=1}^{K} |a_{ik}^{flex} - a_{jk}^{flex}|}{K} \tag{5}$$

Pairwise attitude similarities result in a valued graph, while the original definition of *LS* sets is based on a binary graph. Accordingly, we modified the criterion for a subset *S* of agents to qualify as a cohesive subset such that for each of the nodes in *S*, the weighted sum of the ties to all other nodes in *S* exceeds the weighted sum of ties to nodes outside of *S*. This modification relaxes the original definition also in another way. We do not require that the criterion is met for every subset of *S*, but only for every node of *S*. Without this latter modification the criterion turned out to be far too restrictive for our purposes.

---

[4]The measure presented here is different from the polarization measure used by Flache and Mäs (2008a, 2008b). We tested whether the qualitative results reported in the current article remain identical if we use their measure. No qualitative difference was found. We decided to nevertheless use the new measure, because it is somewhat more sensitive to changes in the distribution of attitudes than the measure employed by Flache and Mäs (2008a, 2008b).

The adapted criterion allows identifying cohesive subsets but it does not yield a numerical value for their degree of cohesiveness. To this end we use a measure that we call *subset segregation* of a subset S, denoted σ^S. This measure is based on the difference of the density of the network of attitude similarities between members of the subset S and the density of the network of attitude similarities between members and non-members of S, as formalized in (6).

$$
\sigma^S = \frac{1}{(N_S - 1)N_S} \sum_i^N \sum_{j \neq i}^N l_i l_j s_{a^{flex}ij}
$$

$$
- \frac{1}{(N - N_S)N_S} \sum_i^N \sum_{j \neq i}^N (l_i(1 - l_j) + l_j(1 - l_i)) s_{a^{flex}ij} \tag{6}
$$

In (6), N_S refers to the number of members of the subset S, N indicates group size, and $l_k = 0$ when agent k is not a member of the focal set S, whereas $l_k = 1$ when k belongs to the focal set. The measure can theoretically vary continuously from -2 to 2 ($-2 \leq \sigma^S \leq 2$). It takes 0 when all agents belong to the set; it takes 2 if the attitudes of the members of the set are maximally similar but at the same time maximally dissimilar from the attitudes of non-members of the set. However, in real graphs it is unlikely that σ^S falls below 0. To align it with our other output measures, we divide this measure by 2, leading to a range of -1 to 1 ($-1 \leq \frac{\sigma^S}{2} \leq 1$). Given that cohesive subsets can vary in size, we additionally weigh σ^S by the extent to which the size of the subset approaches half of the size of the team. Technically, the weighting factor η is obtained as defined in (7). The final measure of polarization (*pol*) is defined in (8). It can vary continuously between 0 and 1 ($0 \leq pol \leq 1$), where 1 indicates maximal polarization as defined above. Within a typical graph there is more than one possible cohesive subset by our criterion; we take the segregation of the most segregated cohesive subset S set as the outcome measure.

$$
\eta = 1 - 2 \frac{|\frac{N}{2} - N_S|}{N} \tag{7}
$$

$$
pol = \eta \sigma^S \tag{8}
$$

Co-variation in fixed and flexible attributes, denoted *cov(fix; flex)*, indicates the extent to which differences in demographic characteristics among team members coincide with differences in attitudes. This is measured as the covariance of the differences in demographic dissimilarities and differences in attitudes across all pairs of agents in the team. Equations (9a) and (9b) define the calculation of the pairwise differences in demographic characteristics and attitudes, respectively. Equation (10) defines how the covariance is obtained.

$$
\Delta_{i,j}^{fix} = \frac{1}{D} \sum_{d=1}^D |a_{id}^{fix} - a_{jd}^{fix}| \tag{9a}
$$

$$\Delta_{i,j}^{flex} = \frac{1}{K} \sum_{k=1}^{K} |a_{ik}^{flex} - a_{jk}^{flex}| \tag{9b}$$

$$cov(fix, flex) = \frac{\sum_{j \neq i}((\Delta_{ij}^{fix} - \overline{\Delta_{ij}^{fix}})(\Delta_{ij}^{flex} - \overline{\Delta_{ij}^{flex}}))}{N(N-1)} \tag{10}$$

Attitude variance, denoted *var*, indicates the 'diversity' in attitudes in a team. This is measured as the average of the standard deviation of each of the K attitudes in a team. The measure can vary continuously from 0 to 1 ($0 \leq var \leq 1$). It takes 0 if there is perfect consensus in a team (i.e. all team members hold maximally similar attitudes). In the case of maximal polarization, as defined here, *var* takes the value 1. However, high values of *var* do not necessarily indicate polarization. For example, *var* can also take comparatively high values in situations in which all agents' attitudes are widely spread across the entire attitude space, without the formation of distinct cohesive subgroups.

Average uncertainty, denoted $\overline{\mu}$, is calculated as the average uncertainty level across all attitudes of all agents. To align it with our other outcome measures, we normalize the measure as described in (11), such that $0 \leq \overline{\mu} \leq 1$.

$$\overline{\mu} = \left(\frac{\sum_{i=1}^{N} \sum_{k=1}^{K} (\mu_{ik,t})}{2NK} \right) \tag{11}$$

4 Results of computational experiments

As indicated in Tables 1 and 2, we varied both faultline strength and initial uncertainty from 0 to 1 in steps of 0.2. This created 36 experimental conditions. Upon initialization, fixed attributes were assigned to agents as indicated in Table 1. Agents' flexible attributes were drawn randomly from a uniform distribution. In the following, we first discuss four exemplary simulation runs in detail. Subsequently, we present the values of the outcome measures, averaged across 50 independent realizations of the simulation experiment for each experimental condition. Each of these realizations had 50,000 iterations.[5]

Figures 3 and 4 depict 'snapshots' of single realizations of the process of attitude and uncertainty formation in four different prototypical conditions. We selected a 2×2 combination of low vs. high faultline strength ($r = 0/1$) and low vs. high initial uncertainty ($\mu_{in} = 0.2/0.8$). In these conditions, faultlines are minimally/maximally strong and should therefore have their weakest/strongest impact on the attitude formation process. However, as discussed below, although initial uncertainty is neither

---

[5]The model was implemented in Delphi 5. The program code can be obtained upon request from the authors.

a – weak faultlines ($r = 0$) and low initial uncertainy ($\mu_{in}= .2$)
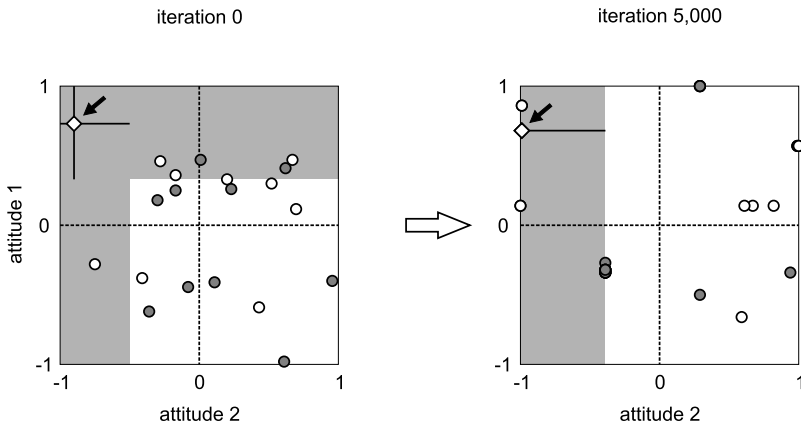
b – weak faultlines ($r = 0$) and high initial uncertainy ($\mu_{in}= .8$)
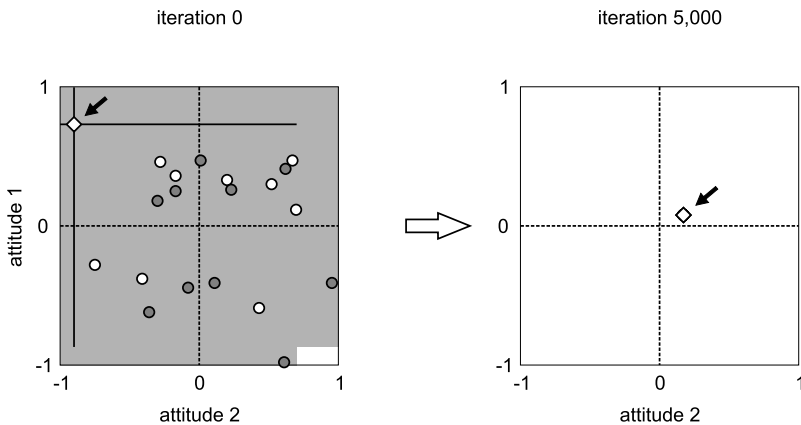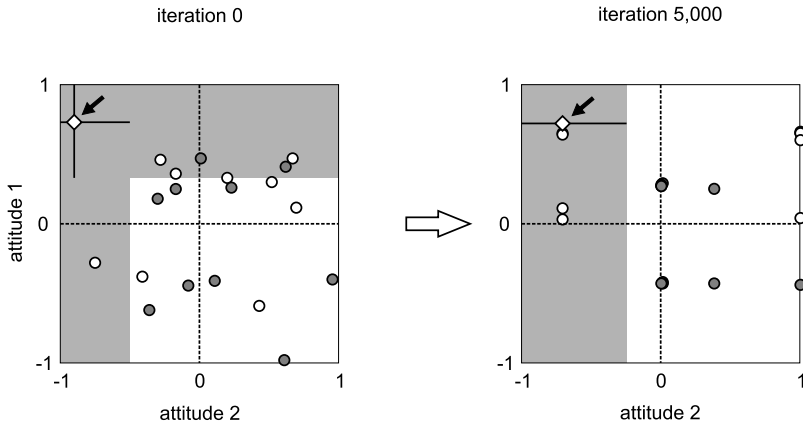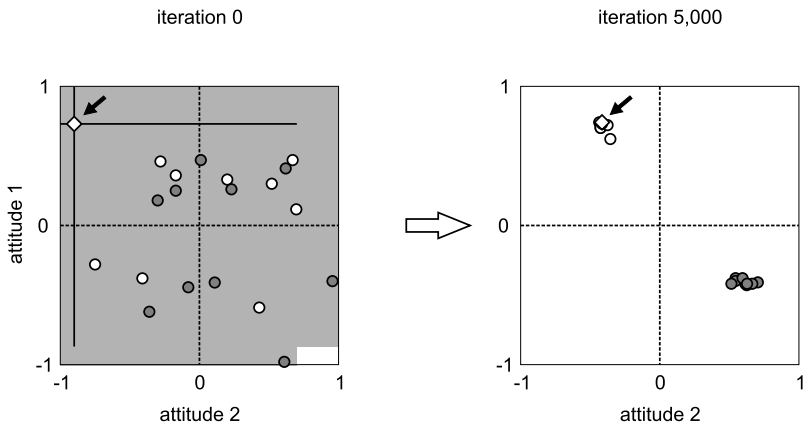

Fig. 3 Exemplary attitude formation processes for two attitudes in teams of 20 agents ($D = 3$, $K = 4$), weak faultlines, and varying initial uncertainty among team members. The shading of *circles/diamonds* represents the value of the first demographic attribute of the respective agent; the *shaded areas* and the error bars show the range of uncertainty for the selected agent (marked as *diamond*) on both issues

minimal nor maximal, it nevertheless has strong effects on the attitude formation process. In both figures each agent is represented by a circle. Each circle's location along the axes represents the values on the first and second of the four attitudes of the respective agent. Note that the initial distribution of attitudes is exactly the same in each of the four examples. We have selected one agent per realization (represented by a diamond) to show her level of uncertainty on both attitudes. This is indicated by gray

a – strong faultline ($r = 1$) and low initial uncertainty ($\mu_{in} = .2$)

b – strong faultline ($r = 1$) and high initial uncertainty ($\mu_{in} = .8$)

Fig. 4 Exemplary attitude formation processes for two attitudes in teams of 20 agents ($D = 3$, $K = 4$), a strong faultline, and varying initial uncertainty among team members. The shading of *circles/diamonds* represents the value of the first demographic attribute of the respective agent; the *shaded areas* and the error bars show the range of uncertainty for the selected agent (marked as *diamond*) on both issues

boxes around the position of the selected agent.[6] The colors of the circles/diamonds indicate the value on the first of the three demographic characteristics of the respective agent (compare Table 1). Note that the coloring of the symbols does not neces-

---

[6]For instance, in part (a) of Fig. 3, the horizontal gray box indicates that the focal agent perceives values between -1 and approximately -0.3 as correct for attitude 1. For attitude 2, the focal agent perceives values between -1 and approximately -0.5 as correct.

sarily reflect similarity relations among agents because dyadic similarity is measured across all demographic characteristics and attitude.

Part (a) of Fig. 3 illustrates the condition in which faultlines are weak and initial uncertainty is low. The figure shows that in this condition the team ends up fragmented into many smaller subgroups whose members hold similar values on one or both of the two attitudes after 5,000 iterations. At the same time, uncertainty does not reach its minimum. Especially for the second of the two attitudes, a considerable level of uncertainty remains as indicated by the gray area around the selected agent. Part (b) of Fig. 3 shows very different dynamics for initially high uncertainty. After 5,000 iterations, the team arrives at perfect consensus on the two issues. At the same time, uncertainty decreases to its minimum for all agents, indicated by the disappearance of the gray boxes in the right subfigure.

The underlying dynamics are in both cases driven by the scarcity of initially 'negative' relations among agents that is due to the low faultline strength. There are potentially many pairs of agents who perceive each other as similar. As a consequence, if initial uncertainty is high (Fig. 3b), positive influence arises in most dyads such that the selected agent is influenced to move toward the other's attitude (positive influence). While in some dyads agents also tend to increase attitude distance, these negative influences are overall outnumbered by positive influences. As a consequence, even dissimilar agents are 'pulled' toward each other because they are positively influenced by the same, or at least similar, sets of peers. In this way, the team quickly arrives at perfect consensus. Once consensus is reached, team members' attitude similarity decreases uncertainty on all attitudes.

If initial uncertainty is low (Fig. 3a), there are only few agents who can actually influence each other in the initial condition. In many dyads agents hold attitudes outside of each others' uncertainty interval in the random initial condition. Yet, some pairs of agents are initially sufficiently close to each other in their attitudes. These agents increase the extent of their agreement with similar others early on in the attitude formation process which, in turn, reduces their uncertainty. The decline of uncertainty countervails further influence from more distant agents and also prevents that disagreement with those agents increases uncertainty.

Next, we explored the effect of initial uncertainty under the assumption of strong faultlines. Part (a) of Fig. 4 shows what happens if faultlines are strong and initial uncertainty is low. After 5,000 iterations, attitudes end up similarly scattered as in the case of weak faultlines and low initial uncertainty. Again, part (b) of Fig. 4 illustrates that the outcome is different if initial uncertainty is high. In this case, the team splits into two opposing (i.e. polarized) subgroups which are internally demographically homogeneous. In this configuration, uncertainty among team members is very low.

Because faultlines are strong in both conditions, there tend to be about as many 'positive' as 'negative' relations among agents at the outset. If initial uncertainty is high (Fig. 4b), this results in either positive or negative influence in most dyads. Demographically similar agents tend to influence each other positively, while demographically dissimilar agents tend to become more dissimilar in their attitudes. Attitude differences then align with the strong demographic faultline that separates the two demographically highly dissimilar subgroups. At the same time, agents' attitudes do not entirely move toward the extreme end of the attitude continuum toward

which their subgroup leans. Furthermore, there is also no 'perfect' agreement within each subgroup. The reason is that disagreement among similar agents is low compared to disagreement among dissimilar agents. This reduces agents' uncertainty and decreases at the same time their strive for further agreement with similar others.

By contrast, if initial uncertainty is low (Fig. 4a), there are at the outset only few agents whose attitudes are sufficiently similar to allow influence. Moreover, in those cases where influence does occur, the outcome of the influence process further reduces agents' uncertainty, countervailing further influence.

The examples presented in Figs. 3 and 4 suggest that the level of initial uncertainty moderates the effect of strong demographic faultlines on attitude dynamics. If initial uncertainty is high, strong faultlines give rise to polarization of demographic subgroups. If uncertainty is low, polarization does not occur. The examples also show that initial uncertainty affects attitude formation if faultlines are weak. If initial uncertainty is high, teams quickly arrive at perfect consensus but consensus does not occur if initial uncertainty is low.

The moderation effect suggested by the exemplary realization is confirmed by the experiments in which we varied both faultline strength and initial uncertainty independently across the full range of conditions shown in Tables 1 and 2. For each experimental condition we computed outcome measures after 50,000 iterations, averaged across 50 independent realizations. Figure 5 reports results for the polarization of demographic subgroups. Part (a) charts the measure of polarization, part (b) displays the alignment of demographic characteristics and attitudes. The figures indicate that both polarization and the alignment measure reach a maximum of $pol = 0.43$ and $cov(fix; flex) = 0.40$, respectively, when both faultline strength and initial uncertainty are maximal ($r = 1/\mu_{in} = 1$). Relative to this level, both measures decrease considerably when the initial uncertainty decreases toward its lowest level ($\mu_{in} = 0$), holding faultline strength maximal ($r = 1$). Polarization drops to 0.20 in this slice of the parameter space and the alignment measure declines to about 0.13. This suggests that in groups with strong demographic faultlines, higher initial uncertainty fosters both subgroup polarization and the alignment of attitude differences with the demographic faultlines.

A different pattern of effects arises when faultlines are weaker. As Fig. 5 illustrates, if faultlines are weakest, both the polarization and the alignment measure vary much less across different levels of initial uncertainty than they do for strong faultlines. For the weakest faultlines, $r = 0$, we found that polarization is consistently at a low level and decreases from 0.16 to 0.03 when μ_{in} is increased from 0 to 1. The alignment measure $cov(fix; flex)$ decreases from 0.03 to 0.01 in the same slice of the parameter space.

The interplay of faultline strength and initial uncertainty also affects the variation of attitudes in teams (measure var). As part (a) of Fig. 6 indicates, for weak faultlines ($r = 0$) var decreases from about 0.29 to about 0.03 with increasing initial uncertainty (i.e. from $\mu_{in} = 0.0$ to $\mu_{in} = 1$). Figure 6 shows further that the result is different if faultlines are strong ($r = 1$). Here, initial uncertainty hardly affects attitude variance. For example, var only decreases from 0.51 to 0.45 between $\mu_{in} = 0$ and $\mu_{in} = 1$. This is because both positive and negative influence are suppressed if initial uncertainty is low. Despite strong faultlines, the resulting attitude distribution is fragmented. But if

Fig. 5 Polarization (*pol*) and co-variation of fixed and flexible attributes (*cov*(*fix*; *flex*)) at varying levels of faultline strength and initial uncertainty; outcome measures are averaged over 50 independent realizations of each of the 36 combinations of faultline strength and initial uncertainty; each realization with 50,000 iterations
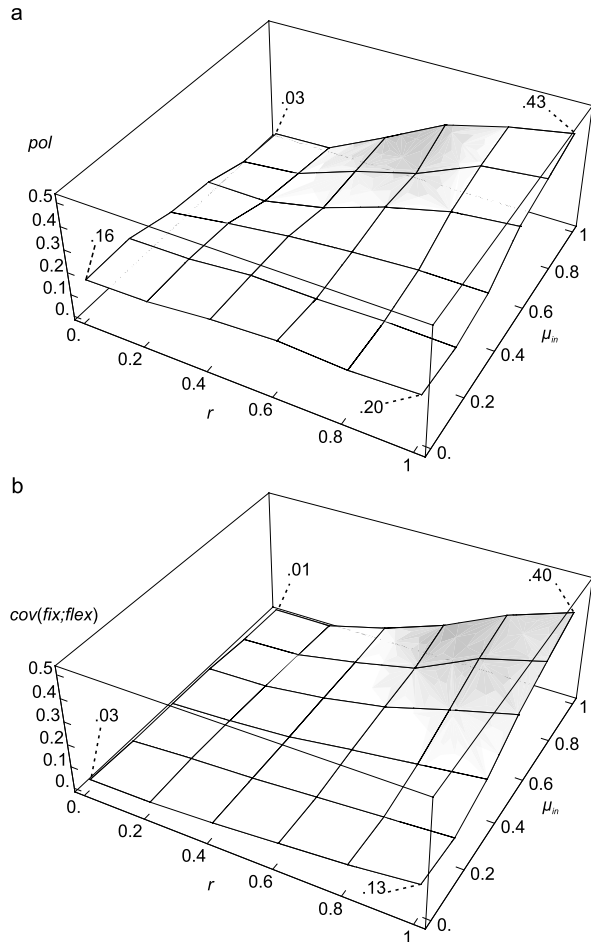
uncertainty is high, agents are open to influence and the team polarizes into opposing subgroups, due to the strong faultline. In both cases there is a comparatively high variation in attitudes in the team. Finally, part (b) of Fig. 6 shows that the mean level of average uncertainty ($\overline{\mu}$) after 50,000 iterations is quite low in all teams. The reason is that the dynamics of the system stabilize when a configuration of attitudes and individual uncertainty levels is reached that brings the further influence process to a hold.

The results show that the averaged outcome measures fall between 0.00 and about 0.50, a range that is relatively small compared to the theoretical boundaries. This reflects that due to the effects of uncertainty our model seldom generates overall 'perfect consensus' or 'perfect polarization' in the sense that all agents hold exactly the same or maximally opposing attitudes. Nevertheless, our results show a clear qualitative pattern: the level of initial uncertainty moderates the effect that faultlines have on the cohesion of teams in terms of subgroup polarization and consensus.

Fig. 6 Attitude variance (*var*) and average uncertainty ($\overline{\mu}$) at varying levels of faultline strength and initial uncertainty; outcome measures are averaged over 50 independent realizations of each of the 36 combinations of faultline strength and initial uncertainty; each realization with 50,000 iterations
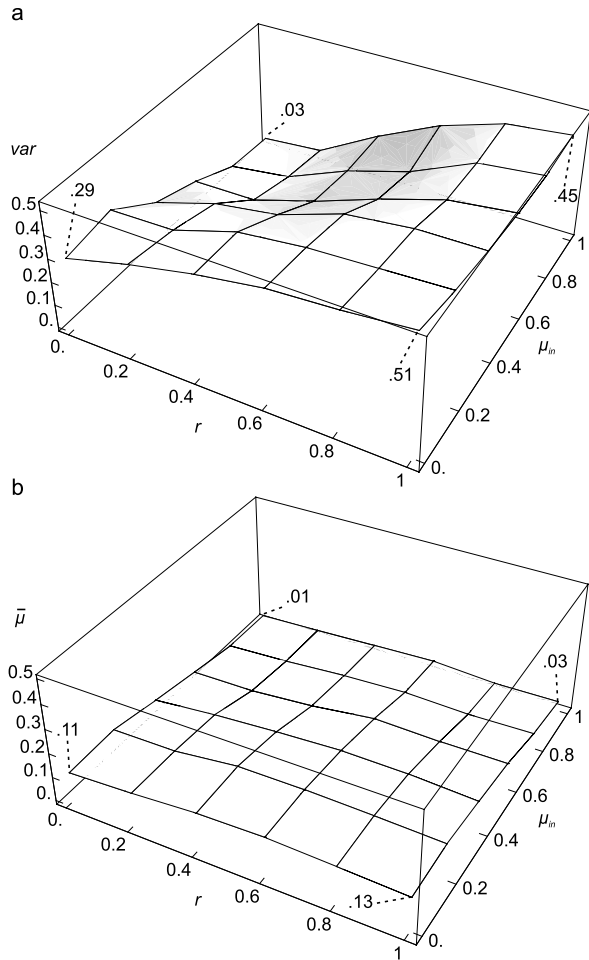
5 Summary and conclusion

Lau and Murnighan's (1998) faultline theory suggests that the stronger demographic faultlines in work teams are, the more likely it is that social identification and influence processes induce polarization of demographic subgroups, alongside with lower cohesion and higher conflict potential in the overall team. Conversely, the weaker faultlines are, the more likely it is that influence processes generate consensus. In this contribution we integrated into faultline theory the insight from social psychology that influence processes depend on the certainty with which individuals hold their attitudes. We argued that this may generate new insights for faultline theory. The level of initial uncertainty among team members may moderate the effect that demographic faultlines have on team cohesion. At the same time, we argued that it is hard to gain solid intuitions about this by informal means, because uncertainty is also endogenous to the attitude formation process.

To test our reasoning, we extended an earlier formalization of faultline theory proposed by Flache and Mäs (2008a, 2008b) toward a formal model of how demographic faultlines affect the interplay of attitude dynamics and the dynamics of uncertainty in a team. Computational experiments with this model generated the expected moderating effect of uncertainty.

This finding also sheds new light on earlier research that assumed effects of uncertainty on the performance of diverse teams (as in Cannella et al. 2008; Carpenter 2002; Carpenter and Fredrickson 2001; Keck 1997). These studies reported mixed effects, with diverse teams sometimes performing better, worse, or not different if uncertainty is high. We found that if diversity in a team is held constant, variations in uncertainty interact with faultline strength in affecting the level of consensus or subgroup polarization. Although consensus and polarization do not directly translate into performance, Lau and Murnighan (1998) suggested that polarization can increase conflict and thereby reduce performance. One might therefore expect that the performance of diverse teams under uncertainty partly depends on the strength of demographic faultlines.

Our research also points to new future directions for theoretical faultline research. We have shown that uncertainty matters, but we have not explored effects of heterogeneity in this agent characteristic. However, research on the effects of extremists on attitude formation processes in groups (e.g. Deffuant et al. 2002; Flache and Torenvlied 2004) suggests that the existence of extremists might importantly affect outcomes of the attitude formation process, if such individuals are less uncertain about their views than moderates. Extremist agents can then lead to polarization in groups by 'pulling' other less extreme and more uncertain agents toward them. This raises the question whether including a connection between extremism and uncertainty in our model would change results. A related intriguing possibility is that the effect of extremists on polarization may be countervailed, if also moderates are sufficiently certain about their attitudes. As Gross et al. (1995) have stressed, the concepts of extremity and uncertainty are analytically distinct. That is, even agents who hold moderate positions might be highly convinced of the correctness of their attitudes. Our model may help to shed light on the conditions under which moderates become less uncertain than extremists and vice versa. For instance, if a moderate position finds much support by other individuals, whereas extreme positions receive less support, moderates might be less uncertain about their evaluation than extremists and therefore might be able to 'pull' extremist toward moderate positions.

Another complication that we have neglected is that the effects of (dis)agreement and (dis)similarity on uncertainty may vary with the salience that a particular issue has for an individual (Holtz and Miller 1985, 2001). The more important an issue is for the identity and self-value of an individual, the more (dis)agreement with relevant others on this issue will affect uncertainty. Hence, teams in which there is low uncertainty about few central issues but high uncertainty about a number of less central issues might experience less polarization than teams in which there is low uncertainty about a large number of less important issues but very high uncertainty about few central issues.

Finally, demographic characteristics and attitudes weigh in our formal model equally in determining agents' similarity, and their relative weight does not change

throughout the entire 'life cycle' of a team. However, Lau and Murnighan (1998) suggested that the relative weights of demographic characteristics and attitudes might change over time; team members may initially base their assessment of others' similarity more on demographic characteristics, while similarity in attitudes becomes more important later on.[7] The reason is that surface demographic characteristics are very accessible whereas individuals need to learn about each others' attitudes during interaction. Flache and Mäs (2008a) already showed that the timing of social interaction can have effects on faultline dynamics. Our findings suggest that timing of who learns when about whose attitude might also be an important issue in relation to uncertainty. For instance, the negative effects of strong faultlines in the case of high uncertainty might be avoided by giving team members only the possibility to learn about the attitudes of those who are similar to them in terms of both demographic characteristics and attitudes. This could decrease their uncertainty and thereby safeguard them from becoming more extreme when they succinctly learn about the attitudes of dissimilar others.

Despite the need for further research, we believe that our theoretical study points to important implications for managerial practice. Our findings suggest that uncertainty in team members' attitudes might be an important determinant for attitude formation processes. Moreover, this determinant interacts with the team's demographic composition. According to our theoretical experiments, managers might benefit from avoiding situations that create high levels of uncertainty among team members if there is a strong demographic faultline within a team. Our integration of faultline theory with theories of social influence suggests that if managers fail to provide objective means for decreasing uncertainty, members of a team with strong demographic faultlines might overly rely on socially provided definitions of the situation, potentially leading to polarization along the faultline between demographic subgroups. If, on the other hand, demographic faultlines are weak, then consensus finding processes might be enhanced if team members 'stick less' to their own attitudes. It appears that in this case actions that make team members somewhat more uncertain in their attitudes (e.g. by providing information that challenges current views within the team) could be beneficial by instigating a reconsideration of the situation by team members.

---

[7]We thank an anonymous reviewer for pointing this out.

References

Abelson RP (1964) Mathematical models of the distribution of attitudes under controversy. In: Frederiksen N, Gulliksen H (eds) Contributions to mathematical psychology. Rinehart Winston, New York, pp 142–160

Blau PM (1977a) Inequality and heterogeneity: A primitive theory of social structure. Free Press, New York

Blau PM (1977b) A macrosociological theory of social structure. Am J Sociol 83(1):26–54. doi:10.1086/226505

Bonabeau E (2002) Agent-based modeling: Methods and techniques for simulating human systems. Proc Natl Acad Sci USA 99(3):7280–7287. doi:10.1073/pnas.082080899

Bordia P, Hobman E, Jones E, Gallois C, Callan VJ (2004) Uncertainty during organizational change: Types, consequences, and management strategies. J Bus Psychol 18(4):507–532. doi:10.1023/B:JOBU.0000028449.99127.f7

Brewer MB (1991) The social self: On being the same and different at the same time. Pers Soc Psychol Bull 17(5):475–482. doi:10.1177/0146167291175001

Cannella AA, Park JH, Lee HU (2008) Top management team functional background diversity and firm performance: Examining the roles of team member colocation and environmental uncertainty. Acad Manag J 51(4):768–784

Carpenter MA (2002) The implications of strategy and social context for the relationship between top management team heterogeneity and firm performance. Strateg Manag J 23(3):275–284. doi:10.1002/smj.226

Carpenter MA, Fredrickson JW (2001) Top management teams, global strategic posture, and the moderating role of uncertainty. Acad Manag J 44(3):533–545. doi:10.2307/3069368

Deffuant G, Amblard F, Weisbuch G, Faure T (2002) How can extremism prevail? A study based on the relative agreement interaction model. J Artif Soc Soc Simul 5(4)

Duncan RB (1972) Characteristics of organizational environments and perceived environmental uncertainty. Adm Sci Q 17(3):313–327. doi:10.2307/2392145

Festinger L (1950) Informal social communication. Psychol Rev 57(5):271–282

Festinger L (1954) A theory of social comparison processes. Hum Relat 7(2):117–140. doi:10.1177/001872675400700202

Festinger L, Schachter S, Back K (1950) Social pressures in informal groups: A study of human factors in housing. Stanford University Press, Stanford

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. doi:10.1007/s10588-008-9019-1

Flache A, Mäs M (2008b) Why do faultlines matter? A computational model of how strong demographic faultlines undermine team cohesion. Simul Model Pract Theory 16(2):175–191. doi:10.1016/j.simpat.2007.11.020

Flache A, Torenvlied R (2004) When will they ever make up their minds? The social structure of unstable decision making. J Math Sociol 28(3):171–196. doi:10.1080/00222500490480257

Gross SR, Holtz R, Miller N (1995) Attitude certainty. In: Petty RE, Krosnick JA (eds) Attitude strength: Antecedents and consequences. Lawrence Erlbaum, Mahwah, pp 215–245

Harrison JR, Lin Z, Carroll GR, Carley KM (2007) Simulation modeling in organizational and management research. Acad Manag Rev 32(4):1229–1245

Hegselmann R, Krause U (2002) Opinion dynamics and bounded confidence models, analysis and simulation. J Artif Soc Soc Simul 5(3)

Hensley V, Duval S (1976) Some perceptual determinants of perceived similarity, liking, and correctness. J Pers Soc Psychol 34(2):159–168. doi:10.1037/0022-3514.34.2.159

Hogg MA, Mullin B (1999) Joining groups to reduce uncertainty: Subjective uncertainty reduction and group identification. In: Abrams D, Hogg MA (eds) Social identity and social cognition. Blackwell, Oxford, pp 249–279

Hogg MA, Turner JC, Davidson B (1990) Polarized norms and social frames of reference: A test of the self-categorization theory of group polarization. Basic Appl Soc Psychol 11(1):77–100. doi:10.1207/s15324834basp1101_6

Holtz R (1997) Length of group membership, assumed similarity, and opinion certainty: The dividend for veteran members. J Appl Soc Psychol 27(6):539–555. doi:10.1111/j.1559-1816.1997.tb00646.x

Holtz R (2003) Intragroup or intergroup attitude projection can increase opinion certainty: Is there classism at college? J Appl Soc Psychol 33(9):1922–1944. doi:10.1111/j.1559-1816.2003.tb02087.x

Holtz R (2004) Group cohesion attitude projection, and opinion certainty: Beyond interaction. Group Dyn 8(2):112–125. doi:10.1037/1089-2699.8.2.112

Holtz R, Miller N (1985) Assumed similarity and opinion certainty. J Pers Soc Psychol 48(4):890–898. doi:10.1037/0022-3514.48.4.890

Holtz R, Miller N (2001) Intergroup competition attitudinal projection, and opinion certainty: Capitalizing on conflict. Group Process Intergroup Relat 4(1):61–73. doi:10.1177/1368430201041005

Holtz R, Nihiser TH (2008) Relative deprivation attitude contrast projection, and opinion certainty. Group Process Intergroup Relat 11(1):89–114. doi:10.1177/1368430207084848

Insko CA, Murashima F, Saiyadain M (1966) Communicator discrepancy stimulus ambiguity, and influence. J Pers 34(2):262–274. doi:10.1111/j.1467-6494.1966.tb01713.x

Jehn KA (1995) A multimethod examination of the benefits and detriments of intragroup conflict. Adm Sci Q 40(2):256–282. doi:10.2307/2393638

Jehn KA, Northcraft GB, Neale MA (1999) Why differences make a difference: A field study of diversity, conflict, and performance in workgroups. Adm Sci Q 44(4):741–763. doi:10.2307/2667054

Jetten J, Hogg MA, Mullin B (2000) Ingroup variability and motivation to reduce subjective uncertainty. Group Dyn 4(2):184–198. doi:10.1037/1089-2699.4.2.184

Keck SL (1997) Top management team structure: Differential effects by environmental context. Organ Sci 8(2):143–156. doi:10.1287/orsc.8.2.143

Lau DC, Murnighan JK (1998) Demographic diversity and faultlines: The compositional dynamics of organizational groups. Acad Manag Rev 23(2):325–340. doi:10.2307/259377

Lau DC, Murnighan JK (2005) Interactions within groups and subgroups: The effects of demographic faultlines. Acad Manag J 48(4):645–659

Macy MW, Willer R (2002) From factors to actors: Computational sociology and agent-based modeling. Annu Rev Sociol 28:143–166. doi:10.1146/annurev.soc.28.110601.141117

McGarty C, Turner JC, Oakes PJ, Haslam SA (1993) The creation of uncertainty in the influence process: The roles of stimulus information and disagreement with similar others. Eur J Soc Psychol 23(1):17–38. doi:10.1002/ejsp.2420230103

McPherson JM, Smith-Lovin L (1987) Homophily in voluntary organizations: Status distance and the composition of face-to-face groups. Am Sociol Rev 52(3):370–379. doi:10.2307/2095356

McPherson M, Smith-Lovin L, Cook JM (2001) Birds of a feather: Homophily in social networks. Annu Rev Sociol 27:415–444. doi:10.1146/annurev.soc.27.1.415

Milliken FJ, Martins LL (1996) Searching for common threads: Understanding the multiple effects of diversity in organizational groups. Acad Manag Rev 21(2):402–433. doi:10.2307/258667

Molleman E (2005) Diversity in demographic characteristics, abilities, and personality traits: Do faultlines affect team functioning? Group Decis Negot 14(3):173–193. doi:10.1007/s10726-005-6490-7

O'Reilly CA III, Caldwell DF, Barnett WP (1989) Work group demography, social integration, and turnover. Adm Sci Q 34(1):21–37

Pomerantz EM, Chaiken S, Tordesillas RS (1995) Attitude strength and resistance processes. J Pers Soc Psychol 69(3):408–419. doi:10.1037/0022-3514.69.3.408

Sherif M, Hovland CI (1961) Social judgment: Assimilation and contrast effects in communication and attitude change. Yale studies in attitude and communication, vol 4. Yale University Press, New Haven

Smith ER, Conrey FR (2007) Agent-based modeling: A new approach for theory building in social psychology. Personal Soc Psychol Rev 11(1):87–104. doi:10.1177/1088868306294789

Smith JR, Hogg MA, Martin R, Terry DJ (2007) Uncertainty and the influence of group norms in the attitude-behaviour relationship. Br J Soc Psychol 46(4):769–792. doi:10.1348/014466606X164439

Smith SM, Fabrigar LR, MacDougall BL, Wiesenthal NL (2008) The role of amount, cognitive elaboration, and structural consistency of attitude-relevant knowledge in the formation of attitude certainty. Eur J Soc Psychol 38(2):280–295. doi:10.1002/ejsp.447

Tajfel H, Turner JC (1986) The social identity theory of intergroup behavior. In: Worchel S, Austin WG (eds) Psychology of intergroup relations, 2nd edn. Nelson-Hall, Chicago, pp 7–24

Thatcher SMB, Jehn KA, Zanutto E (2003) Cracks in diversity research: The effects of diversity faultlines on conflict and performance. Group Decis Negot 12(3):217–241. doi:10.1023/A:1023325406946

Tsui AS, Egan TD, O'Reilly CA (1992) Being different: Relational demography and organizational attachment. Adm Sci Q 37(4):549–579. doi:10.2307/2393472

Turner JC (1985) Social categorization and the self-concept: A social cognitive theory of group behaviour. In: Lawler EJ (ed) Advances in group processes: Theory and research, vol 2. JAI Press, Greenwich, pp 77–122

Turner JC (1991) Social influence. Mapping social psychology. Open University Press, Milton Keynes

van Knippenberg D, Schippers MC (2007) Work group diversity. Annu Rev Psychol 58:515–541. doi:10.1146/annurev.psych.58.110405.085546

van Knippenberg A, Wilke H (1988) Social categorization and attitude change. Eur J Soc Psychol 18(5):395–406. doi:10.1002/ejsp.2420180503

van Knippenberg D, De Dreu CKW, Homan AC (2004) Work group diversity and group performance: An integrative model and research agenda. J Appl Psychol 89(6):1008–1022. doi:10.1037/0021-9010.89.6.1008

Wasserman S, Faust K (1994) Social network analysis: Methods and applications. Structural analysis in the social sciences, vol 8. Cambridge University Press, Cambridge

Wilke H, Young H, Mulders I, de Gilder D (1995) Acceptance of influence in task groups. Soc Psychol Q 58(4):312–320. doi:10.2307/2787131

Williams HM, Meân LJ (2004) Measuring gender composition in work groups: A comparison of existing methods. Organ Res Methods 7(4):456–474. doi:10.1177/1094428104269175

Williams KY, O'Reilly CA (1998) Demography and diversity in organizations: A review of 40 years of research. Res Organ Behav 20:77–140

André Grow received a BA in Social Sciences from the Heinrich-Heine University Düsseldorf, Germany, in 2006 and a MSc in Sociology and Social Research from the University of Utrecht, The Netherlands, in 2009. He is currently a PhD student at the Department of Sociology of the University of Groningen, The Netherlands, and a member of the research school ICS (Interuniversity Center for Social Science Theory and Methodology). His research interests are mainly in social status processes in organizations, team behavior, and computational social and organization theory. For more information see: http://www.rug.nl/staff/a.grow/index.

Andreas Flache is Professor of Sociology at the Department of Sociology of the University of Groningen, The Netherlands, and board member of the research school ICS (Interuniversity Center for Social Science Theory and Methodology). His research addresses mainly cooperation problems and social integration, social networks and learning theory. He applies computational and game theoretical modeling, laboratory experimentation and survey research. Recent publications were included in American Journal of Sociology, Journal of Conflict Resolution, Rationality and Society, Journal of Mathematical Sociology, Journal of Artificial Societies and Social Simulation, and Proceedings of the National Academy of Sciences, among other journals. For more information see: http://www.gmw.rug.nl/~flache.