

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

You can't trust an angry group: asymmetric evaluations of angry and surprised rhetoric affect confidence in trending opinions

Permalink

<https://escholarship.org/uc/item/2dj1d9ds>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

Authors

Richardson, Emory

Keil, Frank

Publication Date

2021

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

You can't trust an angry group: asymmetric evaluations of angry and surprised rhetoric affect confidence in trending opinions

Emory Richardson (emory.richardson@yale.edu)

Frank Keil (frank.keil@yale.edu)

Yale Department of Psychology, 2 Hillhouse Ave. New Haven, CT 06511 USA

Abstract

Communication in groups allows social learners to influence one another and change their beliefs over time. Though some of the same heuristics that guide learners' trust in individual informants can be applied to groups, variation in how individual beliefs are aggregated into a collective judgement can radically alter the accuracy of collective judgement. How do observers evaluate collective judgements? We present two experiments testing the impact of affective signals on observer trust. In each experiment, one faction "converts" group members from an opposing faction, or is converted by them. When the focal faction is *surprised* at the opposing view, observer trust in the focal faction's belief rises or falls as consensus increases or decreases. When the focal faction is *angry*, observer trust falls when consensus decreases, but does not rise even when the "consensus" approaches unanimity. Affective signals in group interactions may help naive learners evaluate collective accuracy.

Keywords: affective signaling, group reasoning, trust in testimony, cooperative learning

Introduction

Widespread misconceptions and misinformation make social learning risky, even as cumulative culture has made humans obligatory social learners: because the skills and knowledge shared by every adult member of a society far exceed what any isolated individual could discover alone, decisions about who to trust are part of our everyday experience. Frequently, however, we place our trust not in individuals, but groups. From vote-counting to grand juries to scientific advancements, collective judgements are often made without any individual having firsthand knowledge of the critical information. While in some cases a collective judgement may be evaluated by collective reputation (e.g., one may trust a vaccine because of the reputation of the manufacturer or the oversight committee), in other cases, one may do better to consider how the beliefs of the individuals involved were aggregated into a collective judgement. Here, we ask how affective signals of individuals in a group discussion — such as anger or surprise at opposing perspectives — impact observers' trust in collective judgements.

Beginning in early childhood, simple heuristics guide learners' trust in individual sources. For example, even infants selectively learn words from sources who have been accurate in the past (Luchkina, Sobel, Morgan, 2018), play more with toys offered by adults who have demonstrated

competence with other toys (Stenberg, 2012), and proactively request labels for novel objects from accurate over inaccurate sources (Bazhydai, Westermann, & Parise, 2020). From an early age, we recognize domains of expertise and rely on domain-experts over non-experts (Lutz & Keil, 2002; Danovitch & Keil, 2004), and we trust kind informants over mean informants (Johnston, Mills, & Landrum, 2015).

We also evaluate sources collectively. Adults and children conform to majority consensus over minority dissent, and majority rule appears to be the default across a variety of contexts, cultures, and even species (Boehm, 1996; Claidière & Whiten, 2012). Majority-based conformity outperforms other social learning strategies in simulation studies (Hastie & Kameda, 2005), and has an evolutionary logic: if selection favors adaptive behaviors, behaviors shared by more agents are more likely to be adaptive. However, in populations of majority-biased conformers, conformity can create bandwagon effects: popular behaviors or beliefs promulgate simply because they are popular. Even in judgements more granular than all-or-nothing imitation, similar "herding" effects can result from people simply adjusting their beliefs to be more similar to their neighbors (Raafat et al., 2009). One way to guard against groupthink is to trust only "independent" firsthand judgements; but strict insistence on firsthand information is implausible in societies with a high degree of specialization and collaborative problem-solving. For example, scientific publications in mature fields regularly have tens or hundreds of authors, none of whom can account for every aspect of their collective work (Hardwig, 1991; Haeussler & Sauermann, 2020). Nevertheless, cutting-edge findings eventually become standard curriculum, taught by instructors many degrees removed from "firsthand" knowledge. Effectively evaluating collective judgement requires some balance of the benefits of collaboration and cumulative culture against the risks of bandwagon effects.

One way to evaluate collective judgement is to consider the process by which individual beliefs are aggregated into a collective response (Dunn, 2019; Goldman, 2014). This can include considering agent-level factors such as the competence or intentions of individual group members (see Harris, Koenig, Corriveau, & Jaswal, 2018, for review). But it can also include group-level factors, such as considering whether opinion changes in a group reflect conscientious error-correction or groupthink, whether the aggregation method is appropriate to the question at hand, or how the aggregation is carried out. For example, even minimal dissent about which of two trajectories will safely land a

probe on a comet might cause concerns about the reliability of the engineers' calculations, and neither majority vote nor taking the average of the conflicting results addresses those concerns; but if one side accepts the other side's calculations, their change of opinion is likely genuine, and may restore confidence. In contrast, a political party whose members toe the party line might cause concern about dogmatism or even political oppression, but both majority vote and "averaging" a middle way are usually considered acceptable forms of collective decision-making. Any decision to trust a collective judgement entails at least implicit endorsement of the process by which individual judgements were aggregated.

Recent work suggests that intuitions about belief aggregation processes develop early. Even as children, people distinguish between contexts in which consensus is more likely to reflect conscientious endorsement or uninformative conformity. For example, in eyewitness memory contexts, children and adults trust a minority with firsthand information over a majority implied to be parroting a single source (Aboody et al., 2019; Mercier & Miton, 2019). Similarly, adults and children believe that polling a large crowd is a better way of inferring the preference of a population than deliberating in a small group — but for mathematical and physical reasoning questions, they expect group deliberation to be more accurate (Richardson & Keil, 2020). Intuitions about the advantages of group deliberation for reasoning tasks are consistent with a large empirical literature demonstrating that groups can outperform not only their average member, but even their best member (Laughlin, 2011; Moshman & Geil, 1998). Yet, despite their potential, groups also regularly *underperform* their potential (Steiner, 1972; Stasser & Stewart, 1992; Kerr & Tindale, 2004), and some work has suggested that people may underestimate both the risks and benefits of some aggregation procedures (Mercier et al., 2020; Mercier et al., 2015; Yousif et al., 2019). Social learners who evaluate collective judgement by the reliability of the belief aggregation process could gain an advantage over learners that ignore group-level processes.

In the present work, we examine how adults respond to one feature of group-level processes: the manner in which individual group members attempt to influence one another. Namely, we ask whether people lose confidence in a faction that attempts to influence opponents by angry shouting, relative to a faction that reacts with surprise to opponent beliefs, but attempts to influence them by simply talking together.

Despite a large literature examining the effects of emotion on individual decision making and interpersonal negotiations, recent reviews note a dearth of research on how emotion influences *group* processes and perceptions of groups (Lerner et al., 2015). Two recent studies compared participants' reactions to violent and non-violent protests as a political tactic. Simpson, Willer, & Feinberg (2018) argue that political violence is likely to reduce support for the violent group's goals because observers see violent groups as "unreasonable", and identify less with those goals as a

result. After reading news accounts about white nationalist and anti-racist protestors that portrayed one or both groups as physically attacking the other, participants' ratings supported the authors' mediation model. Teixeira, Spears, & Yzerbyt (2020) focused instead on norms and the power dynamics between high-status and low-status groups. Results suggested that members of high-status groups in particular see "normative" protests (e.g., petitions, peaceful demonstrations) as effecting change for the low-status group, but see "non-normative protests" (e.g., hacking, traffic-blocking) as simply blaming the high-status group without effecting change for the low-status group. Studying political contexts may make it difficult to separate intuitive reasoning about emotions as an epistemic signal from participants' own political affiliations and their firsthand experience with political protest. Thin-slice evaluations of two orthogonal dimensions, warmth and competence, account for most of the variance in person-perception (Fiske, Cuddy, & Glick, 2007), and tradeoffs between a mean-but-smart and a nice-but-ignorant informant begin to favor the more relevant dimension in moral or "science" questions in early childhood (Danovitch & Keil, 2007; Fusaro & Harris, 2008; Johnston, Mills, & Landrum, 2015). These findings suggest that affective signals could play a significant role in group processes as well.

Evaluations of belief aggregation processes in groups could rely on heuristics similar to warmth-competence evaluations of individuals, but would require reasoning about at least two features unique to group-level processes: the level of consensus among the group members, and how the group members influence one another. For example, the conformity literature suggests higher confidence in greater degrees of consensus, but also a sensitivity to quickly "trending" beliefs, independent of their absolute frequencies (Toelch, Bruce, Meeus, & Reader, 2010). How might affective signals influence observer confidence in changing trends? If affective signals suggest that group members *genuinely* changed their opinion, observer confidence may shift to reflect the final degree of consensus. If affective signals suggest that group members were *forced* to conform, but did not actually change their beliefs, then observer confidence may instead track the degree of consensus prior to the shouting. This suggests an asymmetry between emotions like *Anger* and *Surprise*. While confidence in the accuracy of a faction that expresses *Surprise* at opposing views will *rise* or *fall* as that faction *gains* or *loses* endorsers (reflecting genuine changes in opinion), confidence in the accuracy of a faction that expresses *Anger* at opposing views will *fall* when that faction *loses* endorsers, but it will *not* rise *even when the faction gains* endorsers, because observers will infer that the converts are not endorsing their genuine beliefs. Moreover, if resorting to verbal abuse suggests that a faction is *unable* to persuade opponents by other means, anger may wholly discredit an opinion in the eyes of observers; if this is the case, confidence in the accuracy of the angry faction may *fall* even if the faction *gains* supporters.

We present two experiments exploring the impact of anger on observers' evaluations of collective judgements. In Experiment 1, the *majority* faction expresses *Anger* or *Surprise* at the minority view, resulting in the majority faction either *Gaining* or *Losing* three endorsers (e.g., the majority moves from 6 of 10 agents to 9 of 10 agents or vice versa). Experiment 2 is identical, but the *minority* faction expresses *Anger* or *Surprise* at the majority view (e.g., the minority moves from 4 of 10 agents to 1 of 10 agents or vice versa). In each experiment, we predict that if the change follows a surprised discussion, confidence in the focal faction will rise or fall as they gain or lose endorsers; but if the change follows angry shouting, confidence in the focal faction will fall if they lose endorsers, but will *not* rise even if they gain endorsers.

Experiment 1

Method

Participants. We recruited 60 participants from MTurk for Experiment 1. Each participant saw an Anger trial and a Surprise trial (order counterbalanced). 31 participated in the MajorityGain condition, and 29 in the MajorityLoss condition.

Procedure. Participants were told that each of the 10 students in a classroom setting would answer a test question individually. Then, the students would be allowed to talk together before turning in their test. Participants were told about six topics that the students had studied in science class that year. However, we did not specify for participants which question the students were answering, nor did we reveal what their answers were. Instead, individual student answers were color-coded so that participants knew how many people gave one answer or another. Importantly, this meant that disagreement was not between pre-existing social groups, but opinion-based factions.

Participants were assigned one of two conditions: *MajorityGain* or *MajorityLoss*. In each condition, participants saw a group of 10 students answering a question. A majority of students (6 of 10 or 9 of 10) initially gave one answer, while a minority (4 of 10 or 1 of 10) gave a different answer. Participants were asked to predict which answer was more likely to be accurate. Given the default tendency to trust majority beliefs over minority beliefs, we expected participants to predict that the majority answer was more likely to be accurate. However, the crucial manipulation was the students' emotional response to disagreement. The majority was either "very angry" or "very surprised" that the minority had given a different answer, and "shouted at them" or "talked together with them". As a result of the meeting, three students changed their answer. In the *MajorityGain* condition, three of the four *minority*-opinion students changed their answers after the six majority-opinion students [shouted at / talked with] them; in the *MajorityLoss* condition, three of the nine *majority*-opinion students changed their answers after [shouting at / talking with] the one minority-opinion student. After seeing the opinion change, participants were




Figure 1 (columns read top-to-bottom). Procedure for *Anger* trials in each condition (*MajorityGain* vs *MajorityLose*) of Experiment 1. After being told the possible topics and seeing the initial distribution of beliefs, participants were told, "the students who answered blue were very angry that some students gave a different answer, but after the students who answered green, some of the students changed their answer," and were shown the change. Procedure was identical for Experiment 2, but the *minority* faction was Angry (or Surprised) at the majority faction's response.

asked to make a second prediction about which answer was more likely to be accurate. Participants completed one pre-post Anger trial and one pre-post Surprise trial, rating their relative confidence in the accuracy of the judgements on a 21-point scale.

Results and Discussion

We first analyzed the raw ratings for the Pre-meeting judgments. Past work has suggested that people trust majority over minority judgement, and that confidence increases with the size of the majority. Our data was consistent with this pattern. Participants favored the 6vs4 pre-meeting majority in the *MajorityGain* condition (*Gain3*: $M_{Anger} = 14.48$, $t(30) = 9.12$, $p < .001$, $M_{Surprise} = 13.81$, $t(30) = 7.06$, $p < .001$), and they favored the 9vs1 pre-meeting majority in the *MajorityLoss* condition (*Lose3*: $M_{Anger} = 17.10$, $t(28) = 9.93$, $p < .001$, $M_{Surprise} = 16.93$, $t(28) = 7.33$,

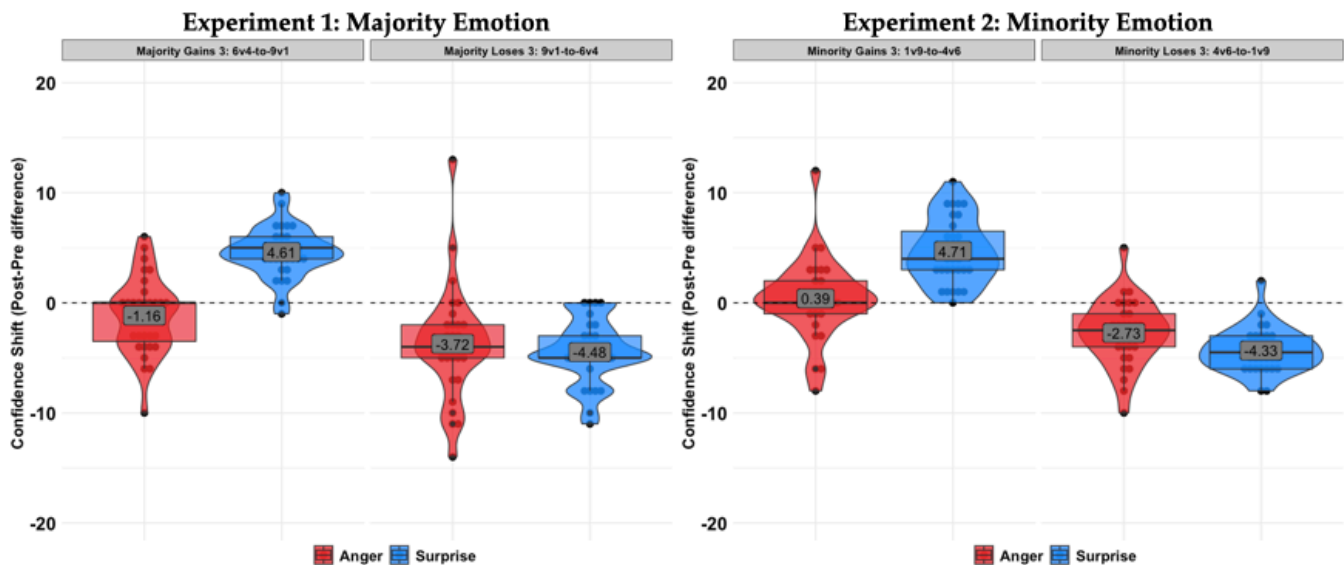


Figure 2. Box and violin plots of confidence shifts in Experiments 1 and 2. Facets show Change condition and pre-post changes in the number of endorsers. For visualization clarity, scores in Exp 2 were multiplied by -1, making positive numbers represent increase in confidence relative to *minority* accuracy. Each point is the difference score for one participant. Grey labels are mean differences.

$p < .001$). Importantly, the pre-meeting judgments within each condition did not differ across *Emotion* ($ps > .2$).

Next, we examined how people's reactions to the discussions differed by the direction of *Change* and *Emotion*. We first computed a difference score by subtracting the Pre-meeting ratings from the Post-meeting ratings for *Anger* and *Surprise* trials, producing negative numbers for a decrease in confidence and positive numbers for an increase in confidence. In each condition, the shift in confidence was consistent with our predictions. When the majority faction was *Surprised* at the minority view, confidence in majority accuracy *rose* significantly when the majority *Gained* three endorsers from the minority view, but *fell* significantly when they *Lost* three endorsers to the minority view (*Surprise*: $M_{\text{Gain}} = 4.61$, $t(30) = 11.10$, $p < .001$; $M_{\text{Loss}} = -4.48$, $t(28) = -8.26$, $p < .001$). In contrast, when the majority faction was *Angry* at the minority view, confidence in majority accuracy fell, regardless of whether the majority *Gained* or *Lost* three endorsers from the minority view — however, the decrease in confidence was only significant in the *Lost* condition (*Anger*: $M_{\text{Gain}} = -1.16$, $t(30) = -1.85$, $p = .074$; $M_{\text{Loss}} = -3.72$, $t(28) = -3.89$, $p < .001$).

The asymmetry between *Angry* and *Surprised* majorities suggests that participants inferred that dissenters who converted to the majority view after being shouted at had not genuinely changed their opinion; moreover, many participants actually reported *less* confidence in the larger post-meeting consensus than the smaller pre-meeting consensus, suggesting that people may have additionally inferred that the angry faction was *unable* to persuade dissenters. However, the backfire effect only approached significance.

Experiment 2

In Experiment 2, we used the same procedure to examine observers' evaluations of *minority* affect. As in Experiment 1, our critical prediction was that confidence in the *Angry* faction would fall when they lost endorsers, but would *not* rise even when they gained endorsers, because participants infer that converts are no longer endorsing their genuine beliefs. However, as “underdogs”, minority factions typically have less power to force majority conformity than vice versa. Hence, it may be less clear that a majority-to-minority conversion is due to the threat of minority anger rather than genuine belief change. A backfire effect may also be weakened for an angry minority: though anger may imply that the angry faction is unable to persuade opponents by other means, the power differential could make it difficult for the minority to convert majority members by *any* means. In short, a plausible alternative is that the power differential could lead participants to excuse the minority's anger altogether, and treat the converts' opinion change as genuine. We can check for an underdog effect that favors minority factions by comparing Experiments 1 and 2: if participants' confidence shifts in Experiment 1 reflected their inference about whether the converts had genuinely changed their belief, then the predicted effect of *Anger* may be weaker in Experiment 2, wherein the *minority* faction is angry, than in Experiment 1.

Method

Participants. We recruited 61 participants from MTurk for Experiment 2. Each participant saw an *Anger* trial and a *Surprise* trial (order counterbalanced). 31 participated in the

MinorityGain condition, and 30 in the MinorityLoss condition.

Procedure. The procedure for Experiment 2 was identical to Experiment 1, except that instead of the *majority* faction expressing anger or surprise at the *minority* view, the *minority* faction expressed Anger or Surprise at the *majority* view. Participants again answered using a 21-point scale.

Results and Discussion

We first analyzed the raw ratings for the Pre-meeting judgments. As in Experiment 1, one-sample *t*-tests for each *Emotion* trial showed that participants favored the 9vs1 pre-meeting majority in the *MinorityGain* condition (Gain3: $M_{\text{Anger}} = 16.87$, $t(30) = 10.60$, $p < .001$, $M_{\text{Surprise}} = 17.06$, $t(30) = 11.94$, $p < .001$), and favored the 6vs4 pre-meeting majority in the *MinorityLoss* condition (Lose3: $M_{\text{Anger}} = 14.10$, $t(29) = 7.12$, $p < .001$, $M_{\text{Surprise}} = 13.77$, $t(29) = 6.13$, $p < .001$). As in Experiment 1, the pre-meeting judgments for each *Emotion* trial did not differ for either condition ($ps > .5$).

Next, we examined how people's reactions to the discussions differed by the direction of *Change* and *Emotion*. As in Experiment 1, we first computed a difference score by subtracting the Pre-meeting ratings from the Post-meeting ratings for Anger and Surprise trials. In order to produce negative numbers for a decrease in confidence and positive numbers for an increase in confidence relative to *minority* accuracy, we then multiplied the difference score by -1. In each condition, the shift in confidence was consistent with our predictions. When the minority faction was *Surprised* at the majority view, confidence in minority accuracy *rose* significantly when the minority *Gained* three endorsers from the majority view, but *fell* significantly when they *Lost* three endorsers to the majority view (*Surprise*: $M_{\text{Gain}} = 4.71$, $t(30) = 8.94$, $p < .001$; $M_{\text{Loss}} = -4.33$, $t(29) = -11.64$, $p < .001$). In contrast, when the minority faction was *Angry* at the majority view, confidence in minority accuracy *fell* significantly when the minority *Lost* three endorsers from the majority view, but did not change significantly when the minority *Gained* three endorsers from the majority view (*Anger*: $M_{\text{Loss}} = -2.73$, $t(29) = -4.96$, $p < .001$, $M_{\text{Gain}} = 0.39$, $t(30) = 0.59$, $p = .56$).

Finally, we compared the *Angry Majority* in Experiment 1 with the *Angry Minority* in Experiment 2; however, evidence of an "underdog effect" was weak. Though participants lost less confidence in the Angry minority than the Angry majority regardless of whether they *Gained* or *Lost* endorsers, the effect was not significant ($M_{\text{Gain}} = -2.40$, $M_{\text{Loss}} = -1.15$, $t(114) = -1.67$, $p = .098$). Regardless of whether the majority shouted at the minority or vice versa, participants appeared to infer that people who converted after being shouted at were not endorsing their genuine beliefs.

General Discussion

Reliance on collective judgement is ubiquitous in modern societies, but the processes by which individual beliefs are aggregated into a collective response are not all equally reliable. While a massive division of cognitive labor may

make it impossible for individuals to evaluate every collective judgement "on the merits", simple heuristics about the reliability of different aggregation processes themselves may allow observers to evaluate collective judgements without access to specialized knowledge. Our experiments suggest that affective signals such as surprise and anger may be one such heuristic.

Recent studies have suggested that emotional language may contribute to the spread of "moralized" content on social media (Brady et al., 2017; cf. Burton, Cruz, & Hahn, 2019); indeed, entire genres of popular fiction are built around themes of anger and vengeance at immoral, ignorant, and incompetent agents, suggesting that people may often see anger as justified in both moral and non-moral contexts. In our experiments, the topics were non-moralized questions about science, and participants had no information about the competing beliefs, allowing a more direct test of the role of emotion per se. Our results complement earlier findings: though observers may use their own prior alliances to justify anger at ideological opponents, our results suggest that anger itself generally engenders doubt about the angry faction's belief, regardless of the apparent popular support for the angry faction.

However, our design also highlights a potential role for agents' *responses* to emotion: though a single person or faction can be angry or surprised for a variety of reasons, observers may interpret responses to that emotion as validating or rejecting the *reasons* for it. For example, observers may interpret steadfastness in one's initial views as a positive cue to accuracy (cf. Christensen, 2009); if a minority remains steadfast in the face of majority anger, observers may shift their support towards the minority faction even if none of the majority faction convert to the minority view.

An even finer distinction might separate the focal faction's *emotions* from their *expression* of that emotion as well as from the other faction's *response*. An angry faction may gain credibility with observers by adopting a "neutral tone" with opponents despite their anger. Of course, adopting a neutral tone may also give a faction's views more legitimacy than they deserve; for instance, naive observers may better recognize absurd claims as such when they are met with open laughter than with a more "neutral" reaction. Moreover, to the extent that our intuitive beliefs about how emotions can act as epistemic signals for observers are accurate, skilled speakers may be able to use them to their advantage in public settings. From a rhetorical standpoint, the "appropriate" reaction when trying to persuade an individual in a private setting may not be the same when trying to persuade the audience in a public setting.

To better understand the epistemic inferences licensed by affective signals, future work should examine a broader range of emotions than those studied here. We contrasted surprise and anger, reasoning that surprise would be the most neutral baseline for comparison to anger short of omitting mention of emotion altogether. Contrasting Surprise with an Omission condition may reveal that surprise licenses its own set of inferences. For example, it's

plausible that omitting mention of emotion altogether would lead observers to treat opinion changes as indicative of general uncertainty in the group instead of error-correction; in this case, participants should moderate their confidence judgments regardless of whether the focal faction gained or lost endorsers.

The ability to reason about emotions is crucial to understanding others' beliefs and desires, and developmental psychologists have long studied children's ability to infer agents' emotions from their beliefs. Only recently has attention turned to the converse problem: inferring agents' beliefs from their emotions (Wu & Schulz, 2018). Our results suggest that affective signals may also influence inferences about others' *knowledge* states. Future work will examine the development of emotion-to-knowledge inferences in children. Changes in the popularity of beliefs may scaffold learners' inferences more effectively than absolute frequencies; simple heuristics like affective signals may help otherwise naive learners distinguish between trends stemming from knowledge-based deference and trends caused by social conformity or Machiavellian maneuvering.

Acknowledgments

We thank the members of the Yale Cognition & Development lab for helpful feedback. This research was supported by NSF grant DRL 1561143 awarded to Frank C. Keil.

References

- Aboody, R., Yousif, S. R., Sheskin, M., & Keil, F. C. (2019). *Says who? Children consider informants' sources when deciding whom to believe*. Cognitive Development Society, Louisville, KY.
- Bazhydai, M., Westermann, G., & Parise, E. (2020). "I don't know but I know who to ask": 12-month-olds actively seek information from knowledgeable adults. *Developmental Science*, 23(5), 10.
- Boehm, C., Antweiler, C., Eibl-Eibesfeldt, I., Kent, S., Knauft, B. M., Mithen, S., Richerson, P. J., & Wilson, D. S. (1996). Emergency Decisions, Cultural-Selection Mechanics, and Group Selection. *Current Anthropology*, 37(5), 763–793.
- Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A., & Van Bavel, J. J. (2017). Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*, 114(28), 7313–7318.
- Burton, J. W., Cruz, N., & Hahn, U. (2019). How Real is Moral Contagion in Online Social Networks? *Proceedings of the Cognitive Science Society*, 7.
- Christensen, D. (2009). Disagreement as Evidence: The Epistemology of Controversy. *Philosophy Compass*, 4(5), 756–767.
- Claidière, N., & Whiten, A. (2012). Integrating the study of conformity and culture in humans and nonhuman animals. *Psychological Bulletin*, 138(1), 126–145.
- Danovitch, J. H., & Keil, F. C. (2004). Should You Ask a Fisherman or a Biologist?: Developmental Shifts in Ways of Clustering Knowledge. *Child Development*, 75(3), 918–931.
- Danovitch, J. H., & Keil, F. C. (2007). Choosing between hearts and minds: Children's understanding of moral advisors. *Cognitive Development*, 22(1), 110–123.
- Dunn, J. (2019). Reliable group belief. *Synthese*.
- Fiske, S. T., Cuddy, A. J. C., & Glick, P. (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in Cognitive Sciences*, 11(2), 77–83.
- Fusaro, M., & Harris, P. L. (2008). Children assess informant reliability using bystanders' non-verbal cues. *Developmental Science*, 11(5), 771–777.
- Goldman, A. I. (2014). Social Process Reliabilism. In J. Lackey (Ed.), *Essays in Collective Epistemology* (pp. 11–41). Oxford University Press.
- Haeussler, C., & Sauermann, H. (2020). Division of labor in collaborative knowledge production: The role of team size and interdisciplinarity. *Research Policy*, 49(6), 103987.
- Hardwig, J. (1991). The Role of Trust in Knowledge. *The Journal of Philosophy*, 88(12), 693–708.
- Harris, P. L., Koenig, M. A., Corriveau, K. H., & Jaswal, V. K. (2018). Cognitive Foundations of Learning from Testimony. *Annual Review of Psychology*, 69(1), 251–273.
- Hastie, R., & Kameda, T. (2005). The Robust Beauty of Majority Rules in Group Decisions. *Psychological Review*, 112(2), 494–508.
- Johnston, A. M., Mills, C. M., & Landrum, A. R. (2015). How do children weigh competence and benevolence when deciding whom to trust? *Cognition*, 144, 76–90.
- Kerr, N. L., & Tindale, R. S. (2004). Group Performance and Decision Making. *Annual Review of Psychology*, 55(1), 623–655.
- Laughlin, P. R. (2011). *Group Problem Solving*. Princeton University Press.
- Lerner, J. S., Li, Y., Valdesolo, P., & Kassam, K. S. (2015). Emotion and Decision Making. *Annual Review of Psychology*, 66(1), 799–823.
- Luchkina, E., Sobel, D. M., & Morgan, J. L. (2018). Eighteen-month-olds selectively generalize words from accurate speakers to novel contexts. *Developmental Science*, 21, 11.

- Lutz, D. J., & Keil, F. C. (2002). Early understanding of the division of cognitive labor. *Child Development*, 73(4), 1073–1084.
- Mercier, H., Dockendorff, M., Majima, Y., Hacquin, A.-S., & Schwartzberg, M. (2020). Intuitions about the epistemic virtues of majority voting. *Thinking & Reasoning*, 1–19.
- Mercier, H., & Miton, H. (2019). Utilizing simple cues to informational dependency. *Evolution and Human Behavior*, 40(3), 301–314.
- Mercier, H., Trouche, E., Yama, H., Heintz, C., & Giroto, V. (2015). Experts and laymen grossly underestimate the benefits of argumentation for reasoning. *Thinking & Reasoning*, 21(3), 341–355.
- Moshman, D., & Geil, M. (1998). Collaborative Reasoning: Evidence for Collective Rationality. *Thinking & Reasoning*, 4(3), 231–248.
- Raafat, R. M., Chater, N., & Frith, C. (2009). Herding in humans. *Trends in Cognitive Sciences*, 13(10), 420–428.
- Richardson, E., & Keil, F. C. (2020). Does informational independence always matter? Children believe small group discussion is more accurate than ten times as many independent informants. *Proceedings of the Cognitive Science Society*, 7.
- Simpson, B., Willer, R., & Feinberg, M. (2018). Does Violent Protest Backfire? Testing a Theory of Public Reactions to Activist Violence. *Socius: Sociological Research for a Dynamic World*, 4, 237802311880318.
- Stasser, G., & Stewart, D. (1992). Discovery of Hidden Profiles by Decision-Making Groups: Solving a Problem Versus Making a Judgment. *Journal of Personality and Social Psychology*, 63(3), 426–434.
- Stenberg, G. (2013). Do 12-Month-Old Infants Trust a Competent Adult? *Infancy*, 18(5), 32.
- Teixeira, C. P., Spears, R., & Yzerbyt, V. Y. (2020). Is Martin Luther King or Malcolm X the more acceptable face of protest? High-status groups' reactions to low-status groups' collective action. *Journal of Personality and Social Psychology*, 118(5), 919–944.
- Toelch, U., Bruce, M. J., Meeus, M. T. H., & Reader, S. M. (2010). Humans copy rapidly increasing choices in a multiarmed bandit problem. *Evolution and Human Behavior*, 31(5), 326–333.
- Yousif, S. R., Aboody, R., & Keil, F. C. (2019). The Illusion of Consensus: A Failure to Distinguish Between True and False Consensus. *Psychological Science*, 30(8), 1195–1204.
- Wu, Y., & Schulz, L. E. (2018). Inferring Beliefs and Desires From Emotional Reactions to Anticipated and Observed Events. *Child Development*, 89(2), 649–662.