

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Stronger evidence isn't always better: A role for social inference in evidence selection and interpretation

Permalink

<https://escholarship.org/uc/item/9f69q5wm>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 40(0)

Authors

Perfors, Amy

Navarro, Denielle J

Shafto, Patrick

Publication Date

2018

Stronger evidence isn't always better: A role for social inference in evidence selection and interpretation

Amy Perfors (amy.perfors@unimelb.edu.au)
University of Melbourne

Danielle J. Navarro (d.navarro@unsw.edu.au)
University of New South Wales

Patrick Shafto (patrick.shafto@gmail.com)
Rutgers University

Abstract

Much of what we know comes from other people, and the quantity of information provided is often constrained by time or space. For a communicator, what information they choose to convey depends not just on the nature of their topic, but also on the social inferences their listeners will make about them based on what they say. For the listener, their interpretation of information given to them depends not just on the information itself, but also on what inferences they make about the bias and motivations of the communicator they received it from. In this paper we explore how and whether these social factors interact with the “true” nature of the information being communicated. We find that stronger evidence does not always lead to stronger conclusions and often leads to increased perceived bias. Communicators, perhaps for this reason and perhaps for others, often modulate the evidence they present to be less unanimous than warranted. This has implications for real-world situations, like communicating about climate change: in such situations, both communicators and listeners behave in what may be individually rational ways, but the end result is that the underlying truth gets distorted.

Keywords: communication; rational inference; evidence selection; climate change; pragmatics

Introduction

In everyday life, people are required to communicate evidence and persuade others in situations where there are strong constraints of time or space. Sometimes the communication limits are externally imposed: journalists have word limits and speakers have time limits. Cognitive limitations impose bottlenecks of their own: a 1000 page report that is never read by anyone cannot be said to constitute effective communication. These constraints present a difficulty; in many situations the body of evidence that exists is quite substantial but only a small subset of it can be communicated to others.

Scenarios such as this, which are common in the real world, create an *evidence selection* problem for the communicator. How should one choose the subset of evidence to give to others? Conversely, how should receivers evaluate evidence when it is given to them? Given that much of what we learn is from other people, these are core theoretical issues in the study of human cognition. Due to the fact that set of the people who collect evidence (e.g., scientists) is not the same as the set who can (and should) benefit from that information, this issue is also of critical practical importance.

The problems facing communicators and receivers are intertwined as well. For instance, one obvious strategy for a

helpful communicator is to select the sample that most closely captures the main statistical information in the true distribution. This is a sensible strategy if the receiver assumes that the evidence has been sampled from the true distribution in a similar way (i.e., helpfully or at least in a random, unbiased manner). Indeed, most theories that focus on how people reason from multiple sources of evidence approach it as a simple estimation problem with this character (Tversky & Kahnemann, 1972; Budescu & Yu, 2006; Trueblood, Kachergis, & Kruschke, 2011). Consistent with this, people often appear to combine multiple sources of evidence roughly by averaging them (Anderson, 1981; Fischer & Harvey, 1999).

However, in the real world, evidence is often *not* generated helpfully or randomly. Rather, it is provided by a communicator with their own capabilities and biases. For instance, suppose you are watching a youtube video you have never watched before. It discusses a new drug that promises to increase sexual prowess. Two sources are quoted predicting that the drug will be 98% and 96% effective respectively. The standard models of evidence aggregation would predict that one ought to conclude that the drug was very effective on the basis of these sources. Yet it also seems plausible that one would instead infer that the reporter was biased since they only selected overwhelmingly favourable reviews. In this social context, because the evidence is *chosen* by an individual, otherwise compelling evidence may not provide strong grounds for believing the event would happen.

The possibility that people reason this way is consistent with a great deal of research demonstrating that making inferences about other people's actions and beliefs is foundational to human cognition (Dennett, 1987; Tomasello, 1999). It plays a key role in human learning (Tomasello, Carpenter, Call, Behne, & Moll, 2005) and communication (Grice, 1975; Frank & Goodman, 2012), and is at least sometimes a rational consequence of the fact that people are choosing the information that they provide to others (Shafto, Goodman, & Frank, 2012). This approach suggests a very different view of the problem of evidence selection and aggregation – one in which learning requires assessing the weight of the evidence *as well as* the biases of the presenter. One goal of this paper is to investigate whether people's behaviour, as both communicators and receivers, is consistent with this possibility.

It is already well-known that people's inferences are shaped by the social context in which they occur. However, most of the existing research focuses on inferences based on the features of the people involved rather than the evidence they are presenting (see Bless, Fiedler, & Strack, 2004, for an overview). In terms of that evidence, there is also ample indication that humans are not very good at monitoring or correcting for sampling processes (e.g., Stasser & Titus, 1985; Kareev, 2000; Fiedler & Juslin, 2006). However, most of this work concerns sampling that is biased due to non-obvious causal mechanisms in the environment (Chater & Oaksford, 2006) rather than the biases of the person providing the information. Given that people are reasonably good at reasoning about others' actions and beliefs in other contexts, might that suggest they would do well in this situation too? Or is the hidden mental state of the communicator enough of a non-obvious causal mechanism to make this difficult?

One thing is clear: when evidence is social in origin the learner should treat inference as a problem of persuasion, at least in part. Viewed as a persuasion problem, it is unclear how people would be expected to behave in the scenario above. Reliance on judgmental heuristics is common, especially among learners who are not highly motivated or who have low cognitive capacity (Kunda, 1990; Petty & Brinol, 2008). But even among those with strong motivation and cognitive resources, different considerations imply different outcomes. On the one hand, a message is more persuasive to an engaged audience if it consists of many strong pieces of evidence (Petty & Cacioppo, 1984; Chaiken, 1980). However, if the sheer extent of the positive evidence seems suspicious, it may also give rise to an impression of bias that may undermine source credibility. The body of evidence showing that people are more influenced by credible sources is extensive (Pornpitakpan, 2004). For instance, people are less persuaded when they know that the speaker has a financial interest in successful persuasion (Hovland & Weiss, 1951) or if they know the speaker's intention to persuade in advance (Hass & Grady, 1975), and they are more persuaded if the speaker does not know they are listening (Walster & Festinger, 1962).

Along similar lines, a suspicious overabundance of supporting evidence may act as a cue about source credibility by turning a superficially persuasive message into an argument that "doth protest too much." Moreover, from the communicator end, people who intuitively believe that this is true may be motivated to *weaken* their own arguments as an attempt to preserve their own credibility. If people behave this way when evidence strongly supports a conclusion, individually sensible behavior could result in situation in which communicators must deliberately distort the truth (lest their reputation suffer), and receivers therefore draw much weaker conclusions than are actually warranted.

This paper presents a preliminary investigation of how people act in this situation, as communicators as well as receivers. In Study 1, people acting as communicators had to provide information about a situation in which the underlying

distribution of evidence is skewed: most experts believe one thing but a single minority dissenter disagrees. In Study 2, receivers were given different patterns of evidence chosen by communicators in Study 1. Do people acting as communicators attempt to select information as veridically as possible, or do they overselect the sole dissenter, perhaps in an effort to appear more unbiased? On the receiver side, are people more likely to conclude that the communicator was biased when the information presented is unanimous? How do prior biases and the specifics of the situation affect these tendencies?

Our findings demonstrate that people communicating to others do not simply select the most mathematically accurate subset of information to provide. Instead, they prefer to include the dissenter even when doing so makes the chosen subset vary markedly from the true distribution. We also show that reasoners make inferences about the biases of the communicator on the basis of the distribution of evidence provided alone. Taken together, these results suggest that social reasoning may play a non-obvious role in how people select evidence for others and interpret evidence themselves.

Study 1: Communicators

Participants. In view of the recent push towards replicability, we ran two identical versions of this study, both times on Amazon Mechanical Turk. Because this was a between-subjects design, with each participant answering only one question, in order to have enough power for all of the sub-questions each version of the study had a large sample size. Of the 587 total participants across both versions, 21 were excluded for failing the attention filter making 566 in total. 354 (62.5%) were male, 522 (92.2%) were from the USA, with 41 (i.e., 7.2%) from India and the rest from a handful of other countries. Ages ranged from 18 to 70 (mean: 32.27). All effects replicated so for ease of presentation we combine the samples and report both together.

Overview. In this between-subjects 2x2 design each person was asked a single question about one scenario after filling out a demographics questionnaire which included a query about their political affiliation and an attentional filter at the end (a multiple-choice question asking what the basic premise of the scenario was). Participants were paid \$0.40USD for the 3-5 minute study. In all conditions, participants were shown a skewed distribution of experts in which all of them strongly endorsed a claim and one dissenter strongly disagreed. In different conditions people saw different numbers of majority experts, each on a different line (e.g., "Person A is $n\%$ confident that X"). The location of each person in the list, including the dissenter, was randomised, as was the assignment of people to the four conditions below.

Cover story manipulation. The two cover stories were designed to elicit different prior biases from the participants. In the POLITICAL CANDIDATE scenario, people were told to imagine being a journalist writing an article about a local political candidate named Hilda Pimlith. The experts in this case corresponded to members of the public whom offered confi-




Figure 1: **Proportion of people including the dissenter as a function of the total number of experts.** Communicators were shown 5, 10, or 15 experts who were all strongly in agreement with each other with the exception of one dissenter. People selected a subset of up to three people to include in a journal article. They showed a strong preference to include the dissenter, even when that dissenter was a tiny minority. Dissenters were just as likely to be included when they were far outnumbered by the other experts.

dence ratings of the probability that they would vote for her. In the CLIMATE CHANGE scenario, participants were still acting as journalists but this time the article was about climate change and the interviews were of scientists who studied it. Regardless of the cover story, people were told their article space was limited and they had to select one, two, or three reviews to include. In both conditions people were also reminded that, as journalists, their job was to accurately report but also that they wanted to be a trusted news source.

Number of experts manipulation. In addition to the cover story, we varied how many interviews each participant was shown. In all cases the majority of experts/interviewees strongly endorsed the claim and one person strongly dissented. The majority experts were each randomly assigned to have between 90% to 96% confidence in the claim: thus, for instance, in the POLITICAL CANDIDATE scenario one interviewee might be 92% confident they were voting for Hilda Pimlith, while in the CLIMATE CHANGE scenario one interviewee might be 92% confident that climate change exists and was caused by humans. The dissenter was always 4% confident. In one condition people saw FIVE total experts; in another they saw TEN and another they saw FIFTEEN.

The varying numbers of experts in each condition means that the average confidence of the experts presented varied by condition as well. With FIVE, the average is 75%; with TEN it is 84%, and with FIFTEEN it is 87%. A set of three experts including the dissenter would thus have an average confidence of 63%, while not including the dissenter would yield an average of 93%. These numbers result in different predictions of what reviewers people should include if their goal is to be as close as possible to the true average they saw. When there are FIVE experts it is closer to the true average to include the dissenter, but when there are TEN or FIFTEEN it is closer to exclude them.




Figure 2: **Role of prior beliefs in communicator behaviour.** Both progressives and conservatives were more likely to include the dissenter than not, and did not vary if the dissenters were more of a minority. Progressives but not conservatives changed their behaviour based on cover story, excluding the dissenter more often when the article was about climate change.

Results

Do people decide to include the dissenter, and does this decision vary as a function of the cover story or the number of experts they saw? As Figure 1 shows, the majority of people in all of the cover story conditions chose the dissenter. Interestingly, there is no difference at all depending on the number of experts they saw: people were just as likely to pick a dissenter if it was one of five experts as if it was one of fifteen (BF of 0.0204 for the effect of number of experts).¹

Figure 2 breaks these results down as a function of cover story and the self-declared political affiliation of the participants. 57.5% of our participants defined themselves as progressive/liberal, while 18.7% affiliated as conservative (the remainder declared themselves as “middle-of-the-road” or declined to state an affiliation and hence are excluded from all analyses comparing progressives and conservatives).

There are several interesting aspects to these results. First, it is clear that there is no effect of the number of experts regardless of political affiliation: in every instance people included the dissenter the majority of the time, and did not vary in their likelihood of including if they were choosing from expert groups of different sizes (for progressives, BF for the model with experts alone is 0.0473; for conservatives, BF is 0.1001). This contradicts what a communicator should do if their only goal was to accurately reflect the mean of true underlying distribution: in that case, when there are 10 or 15 total experts, excluding the dissenter would result in a subset that is closer to the true mean.

Although political affiliation did not affect people’s tendency to include the dissenter at a high rate, it *did* affect how people responded to the cover story. Conservatives did not change their behaviour based on it, including the dissenter at a reasonably high rate regardless (BF for model with cover story alone: 0.2243). By contrast progressives were more

¹Bayes factors (hereafter, BF) reported rely on the Bayesian equivalent of an ANOVA for the difference between these means. The factors were calculated using the BayesFactor package in R (Morey & Rouder, 2014, v.0.9.8). In all cases, analogous frequentist tests were also run with the same qualitative result.

likely to exclude the dissenter if they were writing an article about climate change (BF for model with cover story alone: 4423.05). This may suggest that progressives are aware that the vast majority of scientists believe in climate change (Cook et al., 2013), or may be less likely to assume that people will think they are biased for excluding the dissenter.

So far we have seen that providers tend to include the dissenter even when that dissenter is vastly outnumbered – a strategy that is not in accordance with the behaviour of a person who only cares about accurately reflecting the true distribution. One interpretation for this is that people may also care about something else whose goals are divergent – namely, possibly, the social inferences the receiver of the data will make about them. Perhaps they include the dissenter because excluding them is more likely to make the receiver conclude that they are biased. One way to investigate this is by looking at what inferences receivers make as a function of whether the dissenter is included in the subset or not. We turn to this in the next study.

Study 2: Receivers

Participants. As before, we ran two versions of this 3-5 minute study on Amazon Mechanical Turk, at the same rate of payment. Again all effects replicated so we combine them here. Of the 591 total participants across both versions, 20 were excluded for failing the attention filter, leaving 571 in total. 332 (58.1%) were male, 519 (90.9%) were from the USA, with 48 (8.4%) from India and the rest from other countries. Ages ranged from 18 to 67 (mean: 31.95).

Overview. Again the study involved a fully between-subjects 2x2 design. People were told that they were reading an article that a journalist had written based on interviews of either people considering voting for a POLITICAL CANDIDATE or scientists who study global CLIMATE CHANGE. The article then reports the interviewees (either two or three). This time we varied whether the dissenter was included in the two or three; the non-dissenter always had a majority viewpoint with a confidence between 90% and 96%. For each participant the experts were shown in a random order.

After seeing the reviewers included in the article, people were asked two questions. The *social inference* question asked them what they thought about the author of the article: whether they were impartial; biased; or biased but trying to appear impartial. In order to reduce demand effects, people were also given the option of saying “I have no idea”; these responses were excluded from the analyses.

Participants were also asked a *topic inference* question about what they decided regarding the topic at hand. In the POLITICAL CANDIDATE scenario the question read, “Based on these people, how likely do you think it will be for Hilda Pimlith to get voted into office?” In the CLIMATE CHANGE scenario it read, “Based on these scientists, how likely would you think it is that climate change exists and is caused by humans?” Even though the climate change question asks people to answer based only on the hypothetical article we presented,




Figure 3: **Inferences about bias made by receivers.** Participants in the receiver condition were shown articles either containing a dissenter (present) or not (absent). They were asked what they thought about the author of the article: were they impartial, biased, or trying to appear impartial but not actually (fake)? Those who saw only experts in complete agreement with each other (absent) were far more likely to infer that the communicator was biased.

we did not anticipate that people would be able to divorce that situation from their prior beliefs about climate change. We therefore thought that this question would provide a window into learning how prior beliefs combine with the interpretations people make about socially selected evidence.

Results

As is evident in Figure 3, the social inferences people made were very different depending on whether they saw experts sets with a dissenter or without (BF of 1.75×10^8).² When absent, people are more likely to say that the communicator was biased than give any other alternative interpretation; when present, they are much more even and the modal response is now to conclude that the communicator is impartial.

We can ask more specific questions by breaking down performance by political affiliation and cover story. As before, we include only the 54.9% of our participants identified as progressive and the 17.0% who were conservative. The results, shown in Figure 4, demonstrate that people always make different inferences about bias depending on whether the dissenter is present or not: progressive POLITICAL CANDIDATE (BF = 1.53×10^6); progressive CLIMATE CHANGE (BF = 4.527); and conservative CLIMATE CHANGE (BF = 5.2893388). The least striking difference occurs with the conservatives who read an article about the POLITICAL CANDIDATE (BF = 1.514), but even they were more likely to infer bias when the dissenter was absent.

²This was calculated using the `contingencyTableBF()` function in the BayesFactor package in R, assuming independent multinomial sampling. Chi-squared tests had the same qualitative result.




Figure 4: Inferences about bias made by receivers. Participants were shown articles either containing a dissenter (present) or not (absent). Those who saw only unanimous experts were far more likely to infer that the communicator was biased.

Regardless of condition, then, people always were more likely to infer that the communicator was biased if the dissenter was absent. In addition to this general trend, there were differences based on priors and cover story. Overall, conservatives appeared to have a higher prior belief that the communicator was biased (perhaps because they were more distrustful of the media?). In addition, progressives but not conservatives behaved differently in the CLIMATE CHANGE scenario, being more likely to conclude that the communicator was impartial even when no dissenter was provided.

Having explored how the presence or absence of a dissenter changes people's inferences about bias, we can also ask if it changes the inferences made about the underlying phenomenon of interest. Answers to the *topic inference* question reveal that people's confidence was 82.8% when there was no dissenter but dropped to 69.0% when one was included (BF of 2.32×10^6 that these are different). Given that the underlying average confidence in the TEN and FIFTEEN expert conditions was over 80%, a drop to 69.0% represents a fair degree of miscommunication.

Taken together, these results illustrate the catch-22 situation facing the communicator. If they include the dissenter, people draw inaccurate conclusions, having less confidence in the topic than justified by the range of evidence. But if they do not include the dissenter, they are perceived to be biased. For people in situations in which reputation and trust is essential, this may be an important consideration.

Discussion

Overall, this work suggests that when people are selecting evidence to give to people in a communicative context they are motivated not just by accuracy alone. When the true underlying distribution of evidence is highly skewed, communicators have a strong tendency to include the dissenting opinion, even when that dissenter is vastly outnumbered. Consistent with the idea that this may arise out of a motivation to look impartial, receivers who are presented with evidence showing substantial unanimity are much likelier to infer bias.

The trade-off between reputation and fidelity is especially interesting in light of the fact that many real-life situations have this character. A journalist choosing to report on any issue in which there is substantial but not complete consensus – such as climate change – must risk the possibility of being perceived as biased, if they do not include the small minority of opposing opinions. As our results suggest, including the dissenter means that people who view the evidence emerge with a distorted picture of the underlying reality: they think that the dissenting view is better supported than it actually is. One way to avoid this situation may be to instead provide summary statistics (Lewandowsky, Gignac, & Vaughan, 2013) but given the salience of human-interest stories that focus on a small subset of individuals, this option is not a cure.

Our results provide an interesting contrast with the “weak evidence effect”, which occurs when receiving weak evidence makes people less likely to believe a conclusion than receiving no evidence at all (Fernbach, Darlow, & Sloman, 2011, but see also Harris, Corner, & Hahn, 2013). This effect may arise because the receivers reason that if there were stronger evidence they would have been provided with that. In this paper we find what we might call the “strong evidence effect” in which stronger evidence does not always lead to stronger inferences. Perhaps because strong evidence makes receivers think that the communicator was biased, it causes them to discount the evidence.

That said, we cannot be sure based only on this preliminary study that communicators include the dissenter at such high rates because of their desire to appear unbiased. Multiple other explanations are possible. For instance, since people's confidence decreases when the variance of expert distributions goes up (Budescu, 2006), they may include the dissenter as a way of communicating their unease. Another possibility is that the tendency to include the dissenter is similar to the tendency to overweight small probabilities when they are presented numerically (Hertwig, Barron, Weber, & Erev, 2004). Yet another is that people include the dissenter to communicate something about the ends of the distribution as well as its central tendency. According to any of these explanations, this behaviour on the part of the communicator may have nothing to do with social inference at all. An appealing aspect of the social explanation is that it accounts for the otherwise coincidental fact that receivers *do* interpret the absence of a dissenter to indicate bias. Even so, it is of course possible that receivers make those inferences and communicators don't care or are unable to modify their behaviour to take that into account. We are investigating these other possibilities in ongoing research.

In a sense, our results are consistent with a broad collection of results in the reasoning literature suggesting that people's inferences are inconsistent with normative expectations of how inferences should be drawn from evidence. These normative models assume that evidence is randomly sampled, while the evidence selected by people is clearly not (e.g., Budescu & Yu, 2006; Trueblood et al., 2011). Indeed, a large

literature suggests that people are not unbiased samplers of evidence in general (Stasser & Titus, 1985; Bonawitz et al., 2011; Shafto, Eaves, Navarro, & Perfors, 2012).

Similarly, our results are also consistent with the persuasion literature, although they go beyond existing research on this topic. For the most part, we would expect that a message will be more persuasive if it contains many strong pieces of evidence (Petty & Cacioppo, 1984; Chaiken, 1980). However, when that evidence has been selected by a person whose motivations are unknown, adding it may also induce a “suspicious coincidence” that reduces the effectiveness of the message by undermining the credibility of the source.

Our findings suggest that understanding how people choose evidence – as well as how they evaluate the evidence they have been given – may not simply be a matter of parameter estimation, but may also reflect a sophisticated process of social inference and balancing social goals. Given that people are social creatures, this may seem to be a sensible conclusion; however, this view of evidence and reasoning is one that has not been traditionally been acknowledged by models of evidence selection and aggregation. Future work must explore how robust this finding and to what extent it does actually result from social reasoning rather than alternate causes. In the meantime, our results illustrate one way in which people’s inferences do not simply reflect the objective value of the evidence they are given – stronger evidence is not always better when it comes to human behaviour.

Acknowledgments

This research was supported by ARC grant DP150103280.

References

- Anderson, N. (1981). *Information integration theory*. Hillsdale, NJ: Lawrence Erlbaum.
- Bless, H., Fiedler, K., & Strack, F. (2004). *Social cognition: How individuals construct social reality*. Psychology Press.
- Bonawitz, E. B., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Teaching limits children’s spontaneous exploration and discovery. *Cognition*, 120, 322–330.
- Budescu, D. (2006). Confidence in aggregation of opinions from multiple sources. In K. Fiedler & P. Juslin (Eds.), *Information sampling and adaptive cognition* (p. 327–352). Cambridge University Press.
- Budescu, D., & Yu, H. (2006). To Bayes or not to Bayes? a comparison of two classes of models of information aggregation. *Decision Analysis*, 3, 145–162.
- Chaiken, S. (1980). Heuristic versus systematic information processing and the use of source versus message cues in persuasion. *Journal of Personality and Social Psychology*, 39(5), 752.
- Chater, N., & Oaksford, M. (2006). Mental mechanisms: Speculations on human causal learning and reasoning. In K. Fiedler & P. Juslin (Eds.), *Information sampling and adaptive cognition* (p. 210–238). Cambridge University Press.
- Cook, J., Nuccitelli, D., Green, S., Richardson, M., Winkler, B., Painting, R., ... Skuce, A. (2013). Quantifying the consensus on anthropogenic global warming in the scientific literature. *Environmental Research Letters*, 8(2).
- Dennett, D. C. (1987). *The intentional stance*. Cambridge, MA: Bradford Books/MIT Press.
- Fernbach, P., Darlow, A., & Sloman, S. (2011). When good evidence goes bad: the weak evidence effect in judgment and decision-making. *Cognition*, 119, 459–467.
- Fiedler, K., & Juslin, P. (Eds.). (2006). *Information sampling and adaptive cognition*. Cambridge University Press.
- Fischer, I., & Harvey, N. (1999). Combining forecasts: What information do judges need to outperform the simple average? *International Journal of Forecasting*, 15, 227–246.
- Frank, M., & Goodman, N. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084).
- Grice, H. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and semantics*. Academic Press.
- Harris, A., Corner, A., & Hahn, U. (2013). James is polite and punctual (and useless): A Bayesian formalisation of faint praise. *Thinking and reasoning*.
- Hass, R. G., & Grady, K. (1975). Temporal delay, type of forewarning, and resistance to influence. *Journal of Experimental Social Psychology*, 11(5), 459–469.
- Hertwig, R., Barron, G., Weber, E., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8), 534–539.
- Hovland, C. I., & Weiss, W. (1951). The influence of source credibility on communication effectiveness. *Public Opinion Quarterly*, 15(4), 635–650.
- Kareev, Y. (2000). Seven (indeed, plus or minus two) and the detection of correlations. *Psychological Review*, 107, 397–402.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108, 480–498.
- Lewandowsky, S., Gignac, G., & Vaughan, S. (2013). The pivotal role of perceived scientific consensus in acceptance of science. *Nature Climate Change*, 3, 399–404.
- Petty, R. E., & Brinol, P. (2008). Persuasion: From single to multiple to metacognitive processes. *Perspectives on Psychological Science*, 3(2), 137–147.
- Petty, R. E., & Cacioppo, J. T. (1984). The effects of involvement on responses to argument quantity and quality: Central and peripheral routes to persuasion. *Journal of personality and social psychology*, 46(1), 69.
- Pornpitakpan, C. (2004). The persuasiveness of source credibility: A critical review of five decades’ evidence. *Journal of Applied Social Psychology*, 34(2), 243–281.
- Shafto, P., Eaves, B., Navarro, D. J., & Perfors, A. (2012). Epistemic trust: Modeling children’s reasoning about others’ knowledge and intent. *Developmental Science*, 15, 436–447.
- Shafto, P., Goodman, N. D., & Frank, M. C. (2012). Learning from others: The consequences of psychological reasoning for human learning. *Perspectives on Psychological Science*, 7, 341–351.
- Stasser, G., & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling during discussion. *Journal of Personality and Social Psychology*, 48, 1467–1478.
- Tomasello, M. (1999). *The cultural origins of human cognition*. Cambridge, MA: Harvard University Press.
- Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, 28, 675–691.
- Trueblood, J., Kachergis, G., & Kruschke, J. (2011). A cue imputation Bayesian model of information aggregation. In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*.
- Tversky, A., & Kahnemann, D. (1972). Judgements under uncertainty: Heuristics and biases. *Science*, 185, 1124–1131.
- Walster, E., & Festinger, L. (1962). The effectiveness of “overheard” persuasive communications. *Journal of Abnormal and Social Psychology*, 65(6), 395–402.