Title

Enhancing the Wisdom of Partisan Crowds: Understanding the Role of Sampling Behavior and Social Influence in Bridging Partisan Divides Over Gun Control Policies

Author

Derreumaux, Yrian

Publication Date

2023

UNIVERSITY OF CALIFORNIA
RIVERSIDE

Enhancing the Wisdom of Partisan Crowds: Understanding the Role of Sampling Behavior and Social Influence in Bridging Partisan Divides Over Gun Control Policies

A Dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Psychology

by

Yrian Derreumaux

June 2023

Dissertation Committee:
Dr. Brent L. Hughes, Chairperson
Dr. David Funder
Dr. Daniel Ozer
Dr. Pete Ditto

The Dissertation of Yrian Derreumaux is approved:

Committee Chairperson

University of California, Riverside

ACKNOWLEDGEMENTS

Throughout the past five years, I have been incredibly fortunate to immerse myself in the fascinating world of research, a journey that has been wonderfully enriching and made possible only by the support of my family, friends, peers, and mentors. Once upon a time, my entire reality revolved around the culinary arts, with graduate school merely a distant thought. Yet from that initial spark to the completion of this dissertation, the people dearest to me have steadfastly supported me on this adventure.

First and foremost, I wish to express my profound gratitude to my advisor, Dr. Brent L. Hughes. Brent bravely took a chance on me as one of his first two graduate students, embracing my unconventional path to academia without hesitation. Our rapport was evident from the outset, exemplified by our first conversation where we excitedly discussed potential designs for exploring the role of motivation in shaping information sampling and belief updating. This shared enthusiasm for research, and our mutual love for food, has carried us through countless discussions about broths and regressions alike, making our weekly meetings something I always look forward to. Brent has been an outstanding mentor, always lending a respectful ear, supporting my decisions, and offering encouragement in the face of any outcome. I am deeply grateful for the guidance and patience he has shown me over the years.

My sincere thanks also go to the members of the Social Neuroscience Lab at UCR, past and present. Jake, Eleanor, Julia, Jennifer, Genesis, Arshiya, and Bernice – your camaraderie and kindness have been integral to my academic success. It is rare to find such a warm-hearted and generous group of individuals, let alone in the same lab.

Your support, especially during trying times like the COVID-19 pandemic, has been invaluable in keeping me persevering. Our research would not have been possible without the contributions of numerous research assistants, including the remarkable Kimia Shamsian, who deserves special mention for her exceptional work on a particularly challenging project that culminated in an excellent co-authored paper.

A heartfelt thank you also goes to the UCR Psychology faculty and administration for their time, patience, and expertise. In particular, I am grateful to Dan Ozer and Chandra Reynolds for demystifying experimental design and statistics, transforming my initial trepidation into a newfound favorite subject.

To my friends, I cannot express enough gratitude for your emotional support and generosity in including a financially constrained graduate student in activities like skiing and attending shows. I am forever indebted to you for making this journey not just feasible, but also incredibly enjoyable.

I am deeply grateful to my loving family - my parents, Jean-David and Karen Derreumaux, and my brothers, Sebastien, and Paul-Ivan. Though my research can be nuanced and intricate, you have never wavered in your genuine curiosity and engagement. Your consistent efforts to understand my work, reading my papers, and offering insightful feedback on my presentations have been an invaluable source of support and inspiration. Your dedicated belief in me, even during times of uncertainty, has acted as a beacon of motivation, guiding me through my academic journey. To my parents, Jean-David, and Karen, thank you for your committed faith in my abilities and for fostering an environment that nurtured intellectual curiosity. Your love, encouragement, and the countless sacrifices you have made have all culminated in the person I am today.

Finally, to my amazing wife, Livy, your unwavering support has been the very foundation upon which my success has been built. Your patience during countless late-night work sessions, and your advocacy at every opportunity, have been nothing short of inspiring. Diligently attending numerous practice talks, you have skillfully fine-tuned my delivery, helping me eliminate those dreaded "umms." You have taught me invaluable lessons about the importance of self-care and maintaining a healthy work-life balance, reminding me that professional accomplishments extend far beyond the number and prestige of one's publications. Your wisdom has opened my eyes to the significance of nurturing meaningful relationships and prioritizing self-care in our quest for a fulfilling life. Thank you for being my rock.

ABSTRACT OF THE DISSERTATION

Enhancing the Wisdom of Partisan Crowds: Understanding the Role of Sampling Behavior and Social Influence in Bridging Partisan Divides Over Gun Control Policies

by

Yrian Derreumaux
Doctor of Philosophy, Graduate Program of Psychology
University of California, Riverside, June 2023
Dr. Brent L. Hughes, Chairperson

A core principle of deliberative democracy posits that information exchange enhances the accuracy of group beliefs. However, in the face of unprecedented access to both first-hand empirical information and second-hand estimates from social networks, partisans often disagree on fundamental facts supported by data. In this dissertation, I integrate research on sampling models and motivated reasoning to examine the mechanisms driving partisan disagreements concerning gun control policies and their impact on the wisdom of partisan crowds. Across two studies, partisans learned about the impact of a policy increasing access to guns on subsequent crime rates, with manipulated access to first-hand empirical information, second-hand social estimates, or both. Our findings show that collective error reduces when individuals sample empirical data (Studies 1-2) and further decreases when they consider the average estimates of others (Study 1). The wisdom of crowds is also enhanced when partisans learn from viewing the estimates of fellow partisans, even when they have the choice to decide from which communities to sample (i.e., from Democrats or Republicans). However, collective accuracy was attenuated when partisans sampled social information due to biases in their sampling behavior (e.g., prioritizing in-group members) and the systematic bias of estimates along party lines (Study 2). These findings emphasize the importance of encouraging individuals to diversify their information sources and expand their social networks to include a wider range of perspectives. Furthermore, they reveal the boundary conditions of partisan social influence on the wisdom of crowds, indicating that while social influence can increase collective accuracy to a certain degree, it may also introduce systematic partisan biases that amplify divides, especially when social estimates propagate across networks and become increasingly detached from empirical evidence. In summary, this dissertation highlights the potential benefits and limitations of social influence in shaping collective judgments and offers valuable insights into how individuals gather information and update their beliefs. Ultimately, we hope this research can inform interventions aimed at fostering informed decision-making, bridging ideological divides, and paving the way for collaborative problem-solving in the face of persistent partisan motivations.

Keywords: wisdom of crowds, partisan bias, sampling, collective judgements, gun control

TABLE OF CONTENTS

List of Figures

List of Tables

Introduction

People frequently rely on data to inform their beliefs, their decisions, and to guide their actions, encompassing both life-altering choices (e.g., participating in clinical trials for new drugs) and everyday matters (e.g., planning daily meals). Despite unprecedented access to both first-hand empirical information and the second-hand estimates from peers, partisans often disagree on topics that are supported by data and have scientific consensus, spanning topics such as the impact of gun control measures on crime rate (Smart et al., 2021), beliefs about climate change (Druckman & McGrath, 2019), and the outcome of the 2020 presidential election (Kahn, 2021). This trend poses a series threat to the functioning of a democratic society, which relies on our collective ability to establish common ground based on objective truths (McCoy et al., 2018; Porpora & Sekalala, 2019).

A prevalent explanation for these partisan disagreements posits that political motivations drive individuals to interpret information and draw conclusions that reinforce their existing partisan beliefs, giving rise to biased collective judgements. This phenomenon fits into a long tradition of research on motivated reasoning, which demonstrates how pre-existing beliefs (e.g., favoring one's ingroup) can skew information interpretation towards preferred conclusions rather than accurate ones (Greenberg & Pyszczynski, 1985; Kunda, 1990; Wyer & Frey, 1983). Motivated reasoning has been found to increase with the prominence of group-based identities (Kahan, 2013; Kahan et al., 2017). Consequently, numerous studies have explored interpretive biases resulting from politically motivated reasoning (Cohen, 2003; for meta-analysis see: Ditto et al., 2019; for a review of methods see: Tappin et al., 2020). For instance, research on selective exposure (for meta-analysis, see: Hart et al., 2009) finds that partisans tend to select more congenial relative to uncongenial information (e.g., Albarracín et al., 2005). People not only prefer congenial information but are also motivated skeptics of uncongenial information (disconfirmation bias: Ditto & Lopez, 1992; Taber et al., 2009; Taber & Lodge, 2012). From this perspective, individuals apply varying criteria when assessing information that is inconsistent with their beliefs, as opposed to information that aligns with their views. Consequently, they hold arguments they disagree with to higher standards, demanding more substantial and robust evidence. Collectively, these findings imply that motivated reasoning could be a key contributor to the persistence of partisan disagreements.

Contrary to the motivated reasoning perspective, some cognitive models suggest that biases can originate from processes that are more benign than those associated with group-serving biases. Cognitive psychologists have long used probability theory and statistics as a framework within which to study human statistical inferences (e.g., “intuitive scientists”: Flavell & Ross, 1981; “intuitive psychometricians”: Kunda & Nisbett, 1986; “intuitive statisticians”: Peterson & Beach, 1967). Research in this domain offer alternative explanations for several biases that were previously ascribed to group-based biases (Fiedler, 2000; Meiser & Hewstone, 2006). The main idea here is that, under many conditions, people show a remarkable capacity to intuit different statistical properties in their environment, but they are naïve to the processes that generate them (Fiedler, 2000; Juslin et al., 2007; Lindskog et al., 2013; Nisbett & Wilson, 1977;

Tversky & Kahneman, 1971). As such, people often fail to account for the non-random and unrepresentative nature of their experiences, which can lead to biased evaluations. For example, people tend to discontinue interactions with those who provide negative experiences, regardless of group affiliation. However, they are more likely to have new encounters with ingroup members, which can lead to updating their initial negative perceptions of the ingroup (caused by "a few bad apples"). In contrast, equivalent outgroup impressions may not be updated due to a lack of additional experiences with outgroup members (Denrell, 2005). From this perspective, biases may emerge because of people's sampling environments and their failure to recognize that their experiences were not generated randomly, rather than being driven by group-based motives per se.

Integrating Sampling and Interpretive Sources of Bias

My early graduate work set out to combine these perspectives to understand how partisan biases arise from the interaction of sampling and interpretive sources of bias. In this work now published in the Journal of Personality and Social Psychology, I conducted several studies that empirically test how political motivations influence information processing, from the information people seek out to the interpretations they draw from these experiences (Derreumaux et al., 2022; see Derreumaux, Lindskog, et al., 2023 for a review of sampling biases across both minimal and political groups). In all studies, participants actively sampled numerical information from ingroup and outgroup categories that represented some underlying attribute (e.g., trustworthiness, political knowledge). In each trial, a participant could freely choose to gather information from either the ingroup or outgroup category, and they stopped sampling when they felt confident that they could accurately evaluate each group.

The sampled information was subjected to two manipulations designed to examine how sampling behavior interacts with motivated interpretations to produce biased evaluations. First, we altered the valence of the initial sample to be either overly positive or negative, exploring the impact of first experiences on evaluations and subsequent information sampling. We posited that if most people begin sampling from their own group, this outlier will increase the variability of ingroup relative to outgroup experiences, as their overall ingroup samples will be biased upward or downward. If motivations only affect sampling behavior but not interpretations of the data, differences in participants' ingroup and outgroup evaluations would be similar in magnitude, based on their initial overly positive or negative experiences (i.e., sampling-driven biases). In contrast, if motivations only influence motivated interpretations without impacting sampling, we expected participants to randomly sample from the ingroup and outgroup first, with both positive and negative initial samples leading to more positive ingroup compared to outgroup evaluations. This is because in both situations, people would interpret the information as more favorable than it is and arrive at similarly biased evaluations. Instead, we predicted that biased evaluations would stem from sampling biases, which guide individuals to sample first and most often from the ingroup (Brewer, 2001), generating more variable ingroup experiences (Konovalova & Le Mens, 2020), and that this would interact with interpretive biases (Bergh & Lindskog, 2019), producing asymmetric updating of initial samples based on their congeniality (Hart et al., 2009; Hughes et al., 2017).

The second manipulation involved randomly assigning participants to real group difference conditions, where their ingroup was either (1) better than, (2) worse than, or (3) the same as the outgroup in terms of an evaluative dimension (e.g., trustworthiness). This provided another opportunity to examine the interaction between biased sampled and evaluations. For example, we posited that if participants are not biased in their interpretation of information, and instead base their judgments on the data encountered (data-driven), their evaluations of ingroup and outgroup members should be similarly accurate (i.e., equally likely to report the direction and magnitude of the group difference when the ingroup is doing better or worse). However, we predicted that participants would accurately represent the data they encountered (e.g., (Denrell & Le Mens, 2017; Fiedler, 2000) only when the ingroup was indeed superior. When the ingroup was inferior, we anticipated that participants would fail to recognize the differences and evaluate the ingroup and outgroup similarly (Howard & Rothbart, 1980).

We found that participants tended to sample from their ingroup first and most often, and that this proclivity was driven by ingroup favoritism (e.g., more interest in the ingroup following a positive ingroup expectation) and stronger political convictions. These biases in sampling behavior produced more variable ingroup experiences and predicted more biased evaluations. Moreover, we found that participants dynamically shifted their sampling strategies based on first experiences when the ingroup was worse, obscuring real group differences and allowing participants to maintain plausible deniability (e.g., increased uncertainty) about the inferiority of the ingroup. This uncertainty also manifested in different stopping points and subsequent over-estimations of ingroup averages, suggesting that sampling experiences and uncertainty moderate biased evaluations. Together, our results shed light on the interactive role of sampling and interpretive sources of bias. Specifically, our findings highlight the role of variability in driving biased evaluations as partisans generated and then ‘cherry picked’ from a more varied array of ingroup experiences to maintain favorable ingroup impressions, whereas unfavorable outgroup impressions remain resistant to change due (in part) to fewer and less varied outgroup experiences.

Importantly, this research concentrated on sampling contexts where individuals had first-hand access to empirical data, rather than considering the potential impact of various types of social information and influence on people's beliefs and decisions. In this dissertation, I build on this work by interrogating the ways in which social factors may shape and interact with individual biases in the context of sampling and decision-making processes and its impact on the wisdom of crowds.

Does Social Influence Increase the Wisdom of Crowds?

Past research on the interplay between sampling and motivated reasoning in shaping biased evaluations has primarily focused on biases arising from individuals' sampling behavior of objective, measurable data (e.g., nonpartisan fact-check ratings; Derreumaux et al., 2022). However, the information that people sample in their day-to-day lives often includes data that has been processed and interpreted by other members within their communities. Indeed, in our increasingly interconnected societies, individuals encounter a multitude of judgments that may align and diverge from their own perspectives. A crucial question in the study of human behavior is whether this social influence enhances the wisdom of crowds or not.

Classic research on the "wisdom of crowds"[1] demonstrates that the aggregate estimates of a group frequently surpass the accuracy of any single individual or expert (Galton, 1906; Soll & Larrick, 2009). This principle is anchored in the idea that individual errors and biases tend to cancel each other out when combined, resulting in an estimate that converges toward the true value. The wisdom of crowds operates under the assumptions that: (a) a true answer exists within the population, and (b) despite individual members possibly lacking the right answer, it can still be discerned at the group level by utilizing aggregation techniques to merge the estimates provided by the group members (Centola, 2022). This phenomenon has predicted complex outcomes with astonishing accuracy, covering a wide range of areas such as political and economic forecasting (Wolfers & Zitzewitz, 2004), public policy design (Morgan, 2014), and performance evaluations (Barneron et al., 2019). These findings imply that integrating the judgments of others may be essential for enhancing the accuracy of one's own estimates. In fact, recent efforts have been made to leverage the wisdom of crowds to tackle complex social challenges. For example, to combat misinformation on social media, researchers

[1] The terms "the Wisdom of Crowds" and "Collective Intelligence" are frequently used synonymously. However, the Wisdom of Crowds typically refers to situations in which a group's combined opinions or estimates, which are often diverse and uncorrelated, surpass the accuracy of any single expert or individual. On the other hand, Collective Intelligence typically describes instances where a group of people collectively make more accurate decisions more than any one person, usually through information sharing or collaboration. It is worth noting that we will use the term the wisdom of crowds when examining aggregate estimates and their relation to the ground truth.

launched the "Birdwatch" initiative, which provides Twitter users with collective estimates about the accuracy of a given piece of information (Wojcik et al., 2022).

A fundamental assumption underlying the wisdom of crowds is that individual estimates are neither positively correlated nor systematically biased. However, this assumption is often violated within our interconnected societies (Samimi & Jenatabadi, 2014), especially in political contexts where both the information people seek out and their evaluations of information are influenced by group membership and norms. For instance, individuals tend to seek out information first and most often from their own groups (Bergh & Lindskog, 2019; Derreumaux et al., 2022), and they frequently experience personal pressure to conform to their in-group's perspectives (Brewer, 2001; Feldman, 1984; Hogg & Reid, 2006). This pressure may prompt individuals to align their preferences and opinions with those of their group, further exacerbating biases and correlations in their judgments. The impact of social influence on the wisdom of crowds, however, remains a topic of debate. Some studies have found that social influence enhances the wisdom of crowds (Jayles et al., 2017) even within politically homogeneous social networks (Becker et al., 2019). In contrast, other research indicates that social influence reduces the diversity of individual estimates without reducing collective error, and instead leading to systematic biases, herding, and groupthink (Lorenz et al., 2011; Mavrodiev & Schweitzer, 2021). For instance, a recent study examined whether social influence could explain partisan disagreements over human-caused climate change despite widespread scientific consensus (Guilbeault et al., 2018). The results demonstrated that social influence (viewing estimates from two Democrats and two

Republicans) enhanced the wisdom of crowds, but only when the political affiliations of the estimates were concealed. When participants were informed of the political leanings of the estimates, social learning was diminished, and polarization persisted. These findings raise the possibility that peer-to-peer learning can compromise the accuracy of individual estimates by skewing group estimates away from the true population parameter, especially in contexts with clear group polarization – such as the debates surrounding climate change.

Recent advancements in social network theory offer a parsimonious explanation for the discrepancies observed in the effectiveness of social influence on the wisdom of crowds (for a review, see Centola, 2022). The most critical predictor of whether social influence enhances or diminishes the wisdom of crowds lies in the network structure within which information is exchanged, also known as "networked collective intelligence." Notably, the wisdom of the crowd is preserved when the influence of the most dominant individuals is minimized. This condition is met in decentralized networks, where individuals possess more equal influence, and information and ideas can circulate more freely among people without being filtered or controlled by a central authority. These findings suggest that, under specific circumstances such as in decentralized networks, social influence in partisan networks can enhance collective accuracy.

While significant progress has been made in understanding the conditions under which social learning enhances the wisdom of crowds, several important questions remain unanswered. Firstly, the composition of social networks in prior research has often been artificially prescribed by design. For example, participants are typically randomly assigned to social networks comprising four to six other partisans, and the wisdom of crowds is measured as a reduction in error within and across these small groups. However, people are not usually presented with the average estimates of a fixed number of others, nor are their social circles random. Instead, they have the freedom to choose the configuration of their communities, deciding where they want to gather information and the amount of information they require before forming beliefs. Moreover, there may be pre-existing differences between Democrats and Republicans in terms of their sampling behavior. Our previous work, for instance, reveals that Democrats tend to sample more information than Republicans, yet both Democrats and Republicans arrive at similarly biased evaluations (Derreumaux et al., 2022). When considering the impact of partisan social influence on the wisdom of crowds, sampling behavior may play a larger role in shaping the accuracy of collective judgements compared to when partisans sample first-hand information, which may contribute to the real-world polarization observed in our society. Consequently, it is crucial to investigate whether social influence can enhance the wisdom of crowds in decentralized networks where partisans have the autonomy to determine the makeup and size of their communities, providing a more accurate reflection of everyday information exchange and decision-making processes.

Secondly, the prevailing approach for measuring the impact of social influence involves asking participants for independent estimates on niche statistics from previous years (e.g., unemployment, immigration, military). These independent estimates are then grouped into networks, averaged, and presented to participants, who upon viewing, can update their estimates. The wisdom of crowds is subsequently measured as the average absolute reduction in error after viewing the average estimates of others. Crucially, these questions are often selected because they are difficult to know a priori and cannot be easily answered through quick online searches (Becker et al., 2019). Due to these uneducated guesses, initial estimates tend to vary widely across samples and thus even if individual estimates are systematically biased (e.g., due to partisan biased interpretations), averaging estimates brings people closer to the truth in line with the wisdom of crowds. In other words, these tasks convert entirely uninformed guesses into somewhat less uncertain ones, leading to final estimates that are indeed closer to the actual value, but still, on average, remain considerably distant from the truth. Although numerous factors contribute to the wisdom of crowds, a prominent explanation for the decrease in error of aggregate estimates in partisan social networks is the presence of individuals who possess more knowledge on the topic a priori, making them less likely to update their beliefs following social influence, and consequently pulling the group mean towards the correct answer (Almaatouq et al., 2022). Therefore, an important consideration is whether social influence will increase the wisdom of crowds when individuals can become informed. In other words, is the wisdom of crowds robust to contexts where people have a baseline understanding of the topic, even in a decentralized network? Addressing this question can help us better understand the conditions under which social influence enhances or constrains the wisdom of crowds in real-world settings, thereby informing strategies to harness its potential more effectively.

Finally, in certain learning contexts, people may lack access to first-hand information (e.g., empirical data) and must instead form their beliefs by engaging with others who have already gathered data and developed their own judgments and beliefs. An extensive body of research demonstrates that motivations influence the information people seek out (Taber & Lodge, 2012) or avoid (Ditto & Lopez, 1992), as well as how information is integrated into pre-existing beliefs (Derreumaux et al., 2022; Derreumaux et al., 2023). A critical question therefore is whether social learning enhances the wisdom of crowds compared to aggregating the individual estimates of people with direct access to empirical data. In other words, does averaging the estimates of those who have sampled second-hand information enhance the wisdom of crowds more than taking the average estimate of those who have sampled first-hand empirical information directly? Examining the wisdom of crowds in contexts where individuals either have direct access to objective data or must rely solely on others' opinions offers a more comprehensive understanding of how collective judgements arise in real-world settings.

The Gun Control Debate: Motivated Reasoning of Scientific Evidence

This dissertation focused on the impact of social influence on the wisdom of crowds surrounding a highly contentious issue in the U.S., that of the impact of gun access policies on crime rates. Debates over gun control in the U.S. and disagreements over evidence regarding the effects of gun control policies provide a unique opportunity to study the impact of social influence on the wisdom of crowds within a sampling framework.

For one, gun ownership has emerged as a strong political identifier, encompassing a significant portion of the conservative party (Lacombe, 2019). The increasing prevalence of mass shootings and gun-related deaths in recent years has intensified political pressure from both the left, advocating for gun control legislation, and the right, defending Second Amendment rights (Doherty, 2008; Kleck, 2015). The "shared fear" hypothesis suggests that the deeply rooted division in this debate stems from fear on both sides – fear of firearm casualties due to inadequate gun control and fear of vulnerability due to excessive gun control (Braman & Kahan, 2001; Pierre, 2019). This entrenched fear has solidified unwavering support on both sides, limiting opportunities for dialogue and compromise.

In addition, the limited availability of rigorous scientific studies examining the effects of various gun policies on crime rates exacerbates the polarization (Coates & Pearson-Merkowitz, 2017). Ambiguity in the existing research allows for partisan biases, as interpretations of the scientific evidence often align with party affiliations. For instance, many gun owners are convinced that firearm possession enhances their safety, citing the "more guns, less crime" argument (Kleck & Patterson, 1993; Plassmann & Whitley, 2003) and various reports of successful defensive gun use (Cramer & Burnett, 2012). Critics argue that the effectiveness of defensive gun use is overstated (Wintemute et al., 2010) and maintain that the potential advantages are outweighed by the risks of being threatened or harmed by a firearm (Cook & Ludwig, 2006; Hemenway, 2011). They also reference several case-control studies indicating correlations between gun ownership and increased rates of gun-related homicides or suicides (Anglemyer et al., 2014). This ostensible ambiguity allows individuals to “create their own reality” by selectively gathering information that supports their views while dismissing opposing perspectives.

Lastly, the polarizing nature of gun control obstructs the development of federal legislation aimed at reducing gun violence, despite the staggering number of gun-related deaths in the U.S. (48,830 in 2020 alone, National Center for Health Statistics, 2021). Furthermore, there is a lack of research focused on understanding the psychological factors underlying these disagreements, as well as interventions aimed at promoting constructive engagement and dialogue. Interestingly, despite this divisiveness, there is widespread bipartisan support for "common sense gun reform," such as red flag laws and universal background checks (Pew Research Center., 2021). By examining how individuals weigh and interpret first-hand empirical information and second-hand estimates provided by their peers, we can better understand the cognitive and social factors that contribute to the polarization of opinions despite common goals shared across the aisle. Addressing these challenges requires further investigation into the cognitive mechanisms driving these disagreements and the development of interventions that foster productive engagement and collaboration.

Overview of the Current Research

In this dissertation, I integrate research on sampling and motivated reasoning with research on the wisdom of crowds to elucidate the mechanisms that drive partisan disagreements concerning gun control policies. Sampling models are particularly well-suited for uncovering insights into the way individuals explore information environments.

Understanding information sampling strategies (e.g., where, and how much information people, in addition to how they respond to sampled information), and the extent to which people choose to prioritize social knowledge over their personal experiences, may shed light on the robustness of the wisdom of crowds. Understanding the cognitive mechanisms through which partisan motivations steer sampling and evaluations can pave the way for developing strategies to mitigate these biases, encourage more productive dialogues between opposing parties, and ultimately diminish the influence of partisan divisions on society. Studies 1-2 have been pre-registered and can be accessed through the following links: [Study 1:https://aspredicted.org/M1Q_HTV; Study 2: https://aspredicted.org/CNX_RL8].

Extant research on the impact of social influence on the wisdom of crowds has primarily focused on asking individuals to estimate abstract statistics that are generally unknown, such as historical unemployment rates. Consequently, these estimates are primarily based on conjecture, introducing a level of randomness that, while bringing aggregate estimates marginally closer to the truth, often results in average estimates that remain considerably biased along party lines. It is crucial to consider that people's exchanged information on controversial topics is frequently shaped by their previous experiences and beliefs, which can affect how they learn from others. For instance, individuals who happen to know more about a given topic may be less inclined to rely on the estimates of their peers, which can either benefit or harm group averages depending on the accuracy of their prior knowledge. To address this limitation in past work, we present participants with the range of crime rates in U.S. states for five years before the implementation of a policy that increased access to guns, serving as a benchmark for assessing the change in crime rates following the policy's enactment. By doing so, we emulate real-world scenarios where individuals possess some contextual knowledge to base their expectations and subsequent evaluations, which will likely be guided by their pre-existing beliefs (e.g., that more guns will increase or decrease crime).

In Study 1, we investigate how individuals incorporate the average estimates of a diverse group of individuals after having had unlimited access to government statistics on crime rates (i.e., first-hand empirical information). This approach allows us to determine whether people integrate average social judgements and improve the accuracy of their own estimates when given the opportunity to become thoroughly informed about the issue a priori. We aim to understand the role of prior knowledge and its impact on individuals' ability to leverage the wisdom of crowds, with the goal of enhancing the accuracy of their own estimate.

In Study 2, we explore the propagation of social knowledge over time, examining whether it improves or deteriorates the wisdom of crowds as information spreads across social networks and becomes increasingly distant from the empirical data, akin to the game of telephone. By granting participants the autonomy to determine the composition and makeup of their samples, we assess the resilience of the wisdom of crowds against biases in individuals' sampling and interpretations of social information, as compared to aggregating people’s naïve estimates or their estimates of sampled empirical data. This is particularly relevant because science communication frequently passes through politically charged, homogenous social networks, where partisans may selectively gather and interpret scientific findings that conform to their pre-existing beliefs (e.g., climate change denial; see Druckman & McGrath, 2019). Consequently, these biased interpretations disseminate across social networks and influence others' perceptions and interpretations of the evidence (Watts & Dodds, 2007).

Importantly, across studies, we introduce a financial incentive for participants to enhance the accuracy of their estimates, thus providing a robust test of partisan bias and its impact on the wisdom of crowds. Participants are informed that they will receive bonus payments based on the closeness of their final estimates to the true crime rate, effectively pitting their personal motivations and biases, such as those related to gun control, against the financial rewards associated with accurate estimations. This approach not only increases the overall reliability and validity of the results but also offers a more comprehensive understanding of the interplay between individual biases and the pursuit of objective accuracy. By incorporating financial incentives, these studies aim to reveal the extent to which individuals are willing to adjust their beliefs and estimations in response to the potential for monetary gain, thereby shedding light on the resilience of partisan biases even when confronted with direct personal benefits tied to accuracy.

Study 1: The Influence of Prior Knowledge on the Wisdom of Crowds

Study 1 investigates how individuals incorporate social knowledge (i.e., the average estimate of a group of individuals) after having unlimited access to empirical data. Notably, in this study, social estimates were simulated to be objective indicators, providing an ideal context in which average social estimates are representative sample of the population.

Methods

Participants

Participants (N = 577) were recruited from CloudResearch (see Table 1). In line with our pre-registration, participants were excluded if they did not sample at least 1 time (N = 6) and if they failed the manipulation check (N = 48). Regarding the manipulation check, participants were excluded if they indicated that crime would increase but provided a round 1 estimate that fell below range of 441-45 or indicated that crime would decrease but provided a round 1 estimate that fell above range of 441-451. Regarding outliers, participants were excluded if their estimates exceed the true mean by a factor of 1000 (N = 27). After applying these exclusion criteria, N = 503 participants remained. The sample had a M_{age} of 38.78 (SD = 10.62) and was approved by the Institutional Review Board at a large Southern California University.

We based our sample size on pilot data using *simr* (Green & MacLeod, 2016), testing for the ability to detect a main effect of condition on learning at or above 80% power. A power curve analysis revealed that a sample of at least 500 was necessary to achieve the desired power. We recruited above this minimum to achieve sensitivity for smaller effects and cross-level interactions with individual differences.

Table 1.

Demographic Summary Statistics for Studies 1-2

	Gender	
	Study 1	*Study 2*
Male	42%	33%
Female	57%	65%
Non-Binary	1%	1%
	Race/Ethnicity	
	Study 1	*Study 2*
White	78%	71%
Black	7%	7%
Hispanic	5%	5%
Asian	5%	5%
Mixed Race	6%	12%
	Strength of Political Identity	
	Study 1	*Study 2*
Very Liberal	15%	13%
Liberal	18%	13%
Slightly Liberal	10%	14%
Moderate	8%	8%
Slightly Conservative	15%	13%
Conservative	22%	24%
Very Conservative	11%	8%

Procedure

The experiment consisted of four parts: (1) a cover story, (2) a sampling task, (3) viewing social estimates (experimental condition) versus control (control condition), and (4) personality and demographic questions to test for moderating effects.

Regarding the cover story (part 1), all participants were told that there is an ongoing debate in the U.S. about whether expanding access to guns increases violent crime by increasing the number of people carrying weapons or decreases crime by making it easier for law-abiding citizens to defend themselves from violent criminals. They were then told that government officials had collected crime rate data from counties in U.S. states that expanded access to guns roughly 1 year ago and that we were interested in their perception of whether crime had increased or decreased on average in these counties since expanding access to guns. Importantly, they were told that the data from each county was reported in terms of violent crime rate per 100,000 people, making it possible to directly compare the crime rates between counties to one another. To provide a baseline for changes in crime rates after expanding access to guns, they were told that the average violent crime rate covering the last five 5 years before expanding access to guns ranged from 441 to 451 per 100,000 people[2]. They were then asked to gather crime rate statistics from as many counties as they needed to confidently estimate the new crime rate (for full cover story, see Supplemental Materials).

[2] The mean crime rate, and crime rate range, were based on crime statistics provided by ucr.fbi.gov for 2019. They reflect the average violent crime per 100,000 for five of the top ten U.S. states with the most lenient gun laws.

After viewing the cover story but before beginning the sampling portion of the task, participants provided their prior estimate of the new crime rate, along with a confidence rating for their estimate (round 1 estimate). Prior estimates were collected via a free text entry, with the goal of not anchoring participants to a certain point on the scale. Confidence ratings were collected via a 100-point slider scale, ranging from 0 (no confidence at all) to 100 (total confidence). In addition, participants were asked whether they thought crime had increased, decreased, or stayed the same.

On each trial in the sampling task (part 2), participants selected to sample crime rate statistics, which was represented by a numerical score. They were subsequently given an option to continue sampling, or stop sampling when they felt they could confidently evaluate the new crime rate. The crime rates were 100 integers pulled from a normal distribution with a mean of 446 and standard deviation of 20. Upon sampling, a JavaScript function in Qualtrics retrieved random numbers (with replacement) from the corresponding dataset. The information presented after each sample included an anonymous ID representing the specific county where the crime rate was taken from, along with the crime rate for that county. After viewing an example trial, participants could freely gather as much information as they liked until they felt they had collected enough information to make a judgment.

Once participants decided to stop sampling, they again provided an estimate of the new crime rate via a free entry, along with a confidence rating for their estimate (round 2 estimate), in addition to stating whether they thought crime had increased, decreased, or stayed the same.

To examine whether having access to the estimates of others increases collective accuracy, participants were randomly assigned to a Social or Reflective condition (part 3). In the Social (experimental) condition, participants were told that they had the opportunity to view the average estimate of 10 other mTurk workers, who had previously completed the task in the past, and that they would then have an opportunity to update their own initial estimate. They were then shown their round 2 estimate, along with the average social estimate. In the Reflective condition, participants were told that they had an opportunity to reflect on their prior estimate, and that they would likewise have an opportunity to update their own initial estimate. In the Reflective (control) condition, participants were only shown their round 2 estimate. The social estimates were generated by simulating the sampling behavior of 1,000 people from the empirical data. Specifically, we assumed that each person samples an average of 7 samples (based on prior work; see Derreumaux et al., 2022). Unlike past work, however, the simulations assumed that point-estimates were a one-to-one correspondence with the sampled data (i.e., unbiased estimate), such that each estimate was the mean of the 7 samples. The average estimates were generated by randomly drawing samples of 10 point-estimates (with replacement) and taking the mean. Thus, the social estimates provided to participants were drawn from a more concentrated distribution that closely approximated the true population parameter, with some error.

Following part 3, participants were asked if they wanted to update their prior estimate or not. If participants opted to update their estimate, they again provided an estimate of the crime rate via a free entry, along with a confidence rating for their estimate (round 3 estimate). If participants opted out of updating their estimate, they proceeded to the next phase of the task.

Critically, all participants were financially incentivized to provide accurate estimates of the crime rate. On top of the payment provided for completing the study, participants were also told that they were eligible to win bonus money based on the accuracy of their final estimates. Ultimately, all participants were awarded the same $.50 bonus at the end of the task, regardless of task performance. The financial incentive provides a robust test of the effect of motivations on estimates given the monetary consequence of responding incorrectly. Participants were fully debriefed at the end of the study and were informed that the crime rates they saw were not real and were made up for the purpose of the study.

Figure 1.

Task Diagram for Study 1

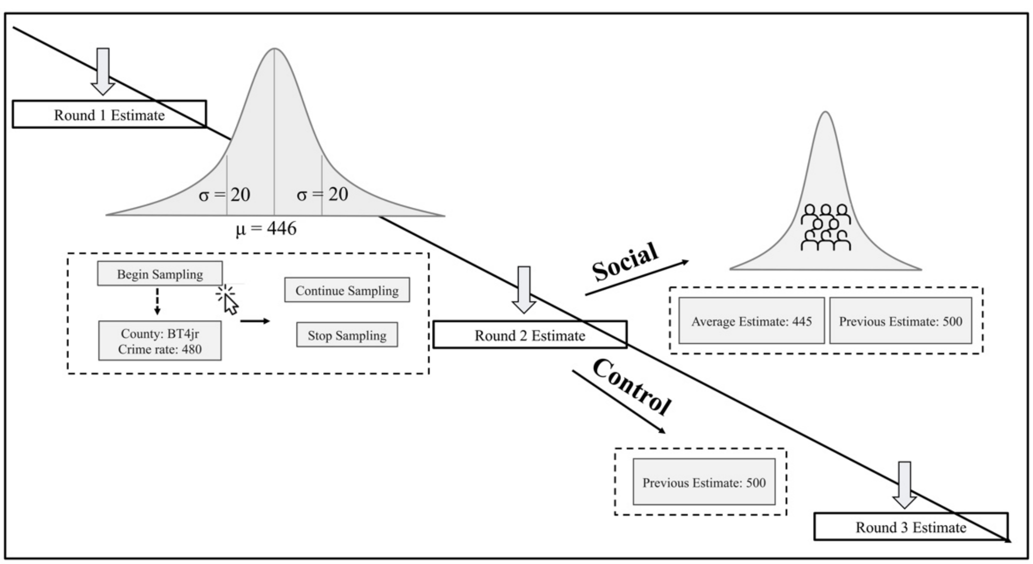

Note. Participants were provided with a baseline for crime rate for the five years prior to the policy change. They could then sample as much empirical information as they needed until they felt confident to estimate the new crime rate. Participants were then randomly assigned to either view the average estimate of ten other CloudResearch workers or reflect on their prior answer.

Self-report measures

Political Identification. Identification with an ideological label was operationalized using a standard 7-point measure of ideological self-description: 1 (*very conservative*), 2 (*conservative*), 3 (*slightly conservative*), 4 (*moderate*), 5 (*slightly liberal*), 6 (*liberal),* and 7 (*very liberal*).

Gun Attitudes Scale (GAS). To assess gun attitudes, participants completed a 9-item general attitude towards gun scale (Tenhundfeld et al., 2020). All items were scored on a 4-point Likert scale ranging from "Strongly Disagree" to "Strongly Agree". The items capture insight into people's beliefs about potential benefits of gun ownership (e.g., "Owning a gun would give me a feeling of independence"). The scale demonstrated excellent reliability in the current sample (a = .95; Omega = .97).

Analysis Plan

Linear mixed models were implemented in R Programming Environment 4.2.3 using lme4 version (Bates et al., 2014) and Satterthwaite approximation degrees of freedom were used for determining p-values in lmerTest (Kuznetsova et al., 2017). Maximal random effects were tested and were removed as needed if unsupported by the data or if the model failed to converge (Barr et al., 2013). Cox proportional hazard models (i.e., frailty models) were implemented using the *coxme* package (Therneau, 2023). Generalized semi-partial R^2 for mixed models were estimated using r2beta (Jaeger et al., 2016).

The Influence of Information Variance and Discordance on Sampling Behavior. To test whether information discordance (i.e., how far sampled information deviated from prior estimates) and variability (i.e., the spread sampled estimates) predicted when participants decided to stop sampling and evaluate, we pre-registered a Cox proportional hazard model. Information discordance was measured as the absolute difference between the running crime rate average that each participant saw across trials and their round 1 estimate. This variable captures changes in how concordant information is on a trial-by-trial basis. Information variability was measured as the running standard deviation of crime rates, capturing the certainty of information over time. Both continuous variables were standardized. The full model included the main effect of information discordance and variance as well as a random intercept for subjects that assumed a normal distribution. This model interrogates whether characteristics in the sampled data influences when participants stop sampling and evaluate.

Connecting Sampling Behavior to Estimates. We pre-registered a linear regression model with the aim of predicting the absolute difference in error before and after sampling the empirical data. Examining visual diagnostics of the residuals, we found that the assumptions of homoscedasticity and normality of residuals were violated (see Figure S1 in Supplemental Materials). As such, we applied a logarithmic transformation to the outcome variable (i.e., absolute difference in error). The full model included an interaction term between a continuous variable representing the standard deviation of the sampled data and the total number of samples each participant gathered.

Do Partisans Tap Into the Wisdom of Crowds? To examine whether people tap into the wisdom of crowds (i.e., the average estimate of others) thereby improving the accuracy of their own estimates, we pre-registered two models. First, we regressed the difference in absolute error between participants' round 2 estimates and the ground truth and participants' round 3 estimates and the ground truth onto a dummy-coded factor representing the condition (with the Reflective condition as the reference group). This measure captures the extent to which people update towards the true crime rate after viewing the average estimate of others compared to when they reflect on their prior estimate.

For the second model, we conducted a Wilcoxon rank-sum test to determine whether the percentage of estimates that moved in the correct direction between rounds was greater in the Social condition compared to the Reflective condition. Specifically, we calculated the percent of estimates that moved toward the true crime rate between rounds 2 and 3. As a final test, we also conducted a chi-square test comparing the proportion of people who updated towards the true crime rate between rounds and conditions. For this model, we coded correct updating as 1 if participants' round 3 estimate was greater than their round 2 estimate, and everything else was coded as 0. This captures whether the proportion of estimates that improved between rounds was greater in the Social condition compared to the Reflective condition.

Association Between Gun Attitudes & the Wisdom of Crowds. We pre-registered several models to test the influence of gun attitudes and political affiliation on the wisdom of crowds[3]. First, to examine whether people's pre-existing gun attitudes influence the accuracy of their estimates, we regressed the difference between people's round 2 estimates and the ground truth onto a continuous measure of gun attitudes, where negative numbers denote unfavorable gun attitudes and positive values denote favorable attitudes.

Next, to test whether gun attitudes influence the extent to which people tap into the wisdom of crowds, we regressed the differences in error between rounds 2 and 3 onto a measure of gun attitudes, as well as a measure capturing the discordance between people's round 2 estimates and the average social estimate they saw. Regarding the discordance measure, negative values indicate a situation where participants underestimated crime relative to the average estimate, while positive values indicate a situation where participants overestimated crime relative to the average estimate. We also measured learning on an individual level by examining whether the proportion of people who update towards the true crime rate differed as a function of gun attitudes using a Wilcoxon rank sum test.

[3] Political affiliation and gun attitudes were positively correlated ($r(501) = -.51$, $p < .0001$) in the current sample. As such, we report models examining gun attitude in-text and report pre-registered political affiliation models in Supplemental Materials (see Table S3-S4).

Results

Sampling Behavior

People Sample Longer in More Variable and Discordant Environments. Our past work found that people's sampling behavior (e.g., sampling first and most from their ingroup) generated greater variability in ingroup experiences, which predicted when people decided to stop sampling, and lead to more biased evaluations (Derreumaux et al., 2022). We first sought to replicate this effect in a context where people sample empirical data about crime rates by examining whether characteristics of the sampled data were associated with when people decide to stop sampling and evaluate.

Replicating past work and in line with our pre-registered hypothesis, we found that both information variance (HR = .73, 95% *CI* [.58, .92], $z = -2.65$, $p = .008$) and information discordance (HR = 0.72, 95% *CI* [.52, .99], $z = -2.05$, $p = .041$) were significantly associated with the hazard of stopping. Specifically, a one-unit increase in information variance was associated with a 26.6% decrease in the likelihood of stopping while holding information discordance constant, and a one-unit increase in information discordance was associated with a 28.0% decrease in the likelihood of stopping while holding information variance constant. We did not observe a significant information variance by discordance interaction (HR = 1.11, $p = .19$). These findings suggest that people are more likely to continue sampling to the extent that their experiences are more (compared to less) variable and more (compared to less) discordant.

Evaluations

Connecting Sampling to Downstream Evaluations. We found that partisans sampled longer in more variable and discordant environments. This could be attributed to people responding to increased uncertainty by engaging in more extensive sampling, resulting in more accurate estimates. Alternatively, people may survive longer by selectively seeking congenial information in uncongenial environments. To test these competing hypotheses, we examined whether information variance and the number of trials participants sampled predicted collective accuracy.

Results revealed a weak yet significant variance by sample interaction (b = .002, 95% CI [.00004, .005], SE = .001, t = 1.99, p = .046, see Figure 2), indicating that individuals who sample longer in more variable environments tend to exhibit more biased evaluations rather than less biased ones. In line with previous research, these findings indicate that people who sample more information in variable environments tend to develop more biased evaluations, rather than becoming less biased as one might expect, highlighting the complex relationship between environmental factors and the development of partisan biases.

Figure 2.

Trials and Variance Associated With Biased Estimates

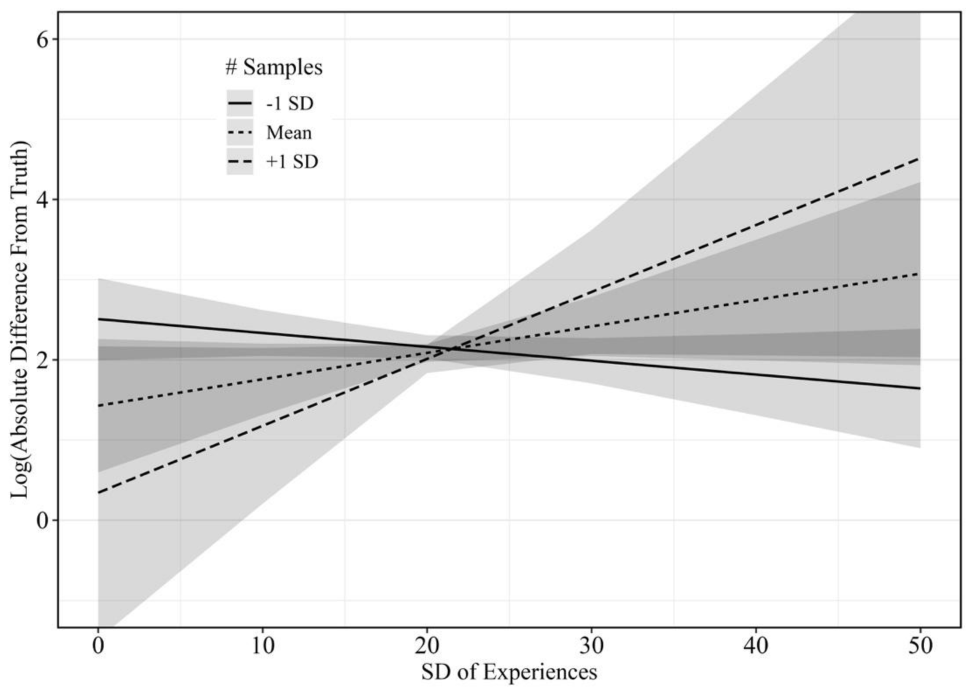

Note. Higher standard deviation of experiences (i.e., sampled data) and number of samples was associated with greater deviations from truth. The y-axis represents log(absolute error). Error bars denote 89% CI around the mean.

Partisans Tap Into the Wisdom of Crowds. One aim of Study 1 was to explore the process through which individuals reconcile differences between their own recently informed estimates and the average estimates of others. To this end, we investigated whether participants did indeed learn from sampling empirical data, and subsequently assessed whether partisans tap into the wisdom of crowds, thereby increasing collective accuracy. A Wilcoxon rank sum test revealed a significant difference in absolute error before and after participants sampled (W = 203,836, $p < .001$), which corresponds to an

80% median improvement in accuracy following sampling (see Figure 3, Table 2). This finding highlights the capacity for individuals to learn from their “hands-on” experience, even surrounding a the highly contentious topic such as gun control.

Next, we examined whether people tap into the wisdom of crowds, thereby increasing the accuracy of their own initial estimates and we observed that collective accuracy was greater for participants in the Social condition compared to Reflective condition ($\beta = -.18$, 95% $CI = [-.36, -.01]$, $SE = .09$, $t = -2.11$, $p = .035$, see Figure 3). This suggests that people can be wise to the wisdom of crowds, integrating social knowledge to improve the accuracy of their own judgements.

Finally, we also tested whether the proportion of people who updated towards the true crime rate was greater in the Social compared to Reflective conditions. A Wilcoxon rank sum test with continuity correction revealed that a higher proportion of individuals in the Social condition updated in towards the true crime rate compared to the Reflective condition ($W = 182$, $p < .0001$), providing converging evidence that people tap into the wisdom of crowds.

Figure 3.

Normalized, Truth-Centered Mean at Each Round

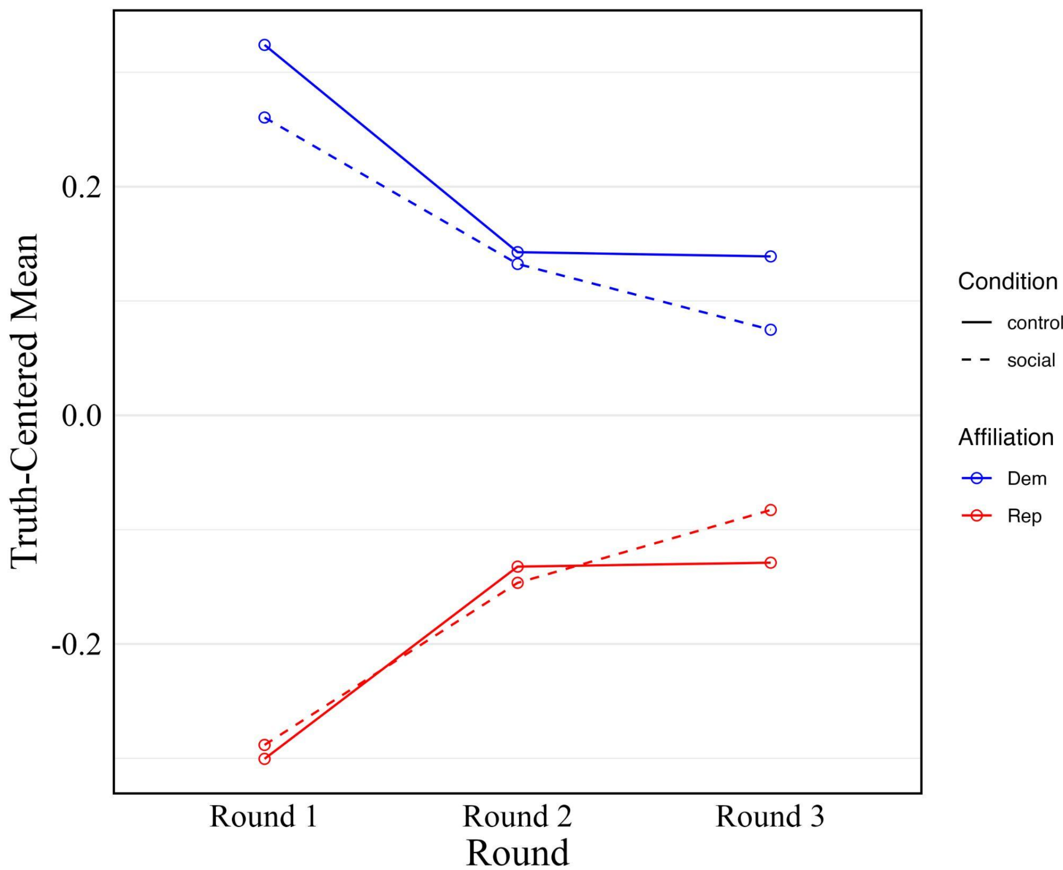

Note. The value for each data point is obtained by calculating the arithmetic difference between the mean belief and the true value at each round and then averaging this value across participants for each political party. The Control condition is denoted via solid lines, whereas the Social condition is denoted via dotted line. Red lines indicate Republican responses whereas blue lines indicate Democrat responses.

Table 2.

Descriptive Statistics of Crime Rate Estimates

		Study 1		Study 2	
		Mean	*SD*	*Mean*	*SD*
Democrat	Round 1	524.66	143.80	457.25	197.28
	Round 2	451.45	42.05	452.63	38.23
Republican	Round 1	431.57	154.61	388.77	185.69
	Round 2	437.59	58.77	444.72	40.68

Note. Mean and standard deviation of crime rate estimates for Democrats and Republicans at Rounds 1-2 for Studies 1-2. The true crime rate was 446.

Extreme Gun Attitudes Undermine the Wisdom of Crowds. The findings thus far present a promising picture of partisan information processing as collective judgements were more accurate after participants sampled empirical data. Furthermore, partisans were wise to the wisdom of crowds, integrating the average estimates of others. One important consideration, however, is whether pre-existing gun attitudes impede this learning process.

We first tested whether partisans were biased in their estimates of the empirical data, arriving at conclusions that align with their pre-existing beliefs, thereby undermining the wisdom of crowds. To that end, we regressed the difference between participant's estimates and the true crime onto a measure of gun attitudes. This model revealed a significant negative effect of gun attitudes on deviation from truth (β = -.13, 95 % CI = [-.21 -.04], SE = .04, t = -3.01, p = .002, see Figure 4A)[4], suggesting that partisans with unfavorable gun attitudes overestimated the true crime rate, while those with favorable gun attitudes underestimated it.

Subsequently, we investigated whether gun attitudes were associated with the extent to which people tapped into the wisdom of crowds based on the discordance between the average social estimate people saw and their own round 2 estimate. To that end, we regressed the difference between participant's estimate and the true crime rate before and after they viewed the average social estimate onto a measure of gun attitudes,

[4] We observed a similar effect of political affiliation, demonstrating that Democrats overestimated the true crime rate, whereas republican underestimated it (β = -.26, 95 % CI = [-.44 -.09], SE = .09, t = 2.98, p = .003).

as well as a variable representing information discordance. Regarding the latter, this was calculated by taking the difference between participants' round 2 estimates and the average social estimate they saw, where positive values indicate participants underestimating crime relative to the social estimate, and negative values indicate participants overestimating crime relative to the social estimates. The model revealed a significant gun attitude by information discordance interaction (β = -.37, 95 % *CI* = [-.46 -.28], *SE* = .04, *t* = -8.16, *p* < .0001, see Figure 4B). We interrogated this interaction with a simple slope analysis, revealing a negative slope of gun attitudes on learning when participants underestimated crime relative to the social estimate (-1 *SD* below the mean of Info Discordance: β = -.49, *SE* = .06, *t* = -7.82, *p* < .001), and a positive slope of gun attitudes on learning when participants overestimated crime relative to the social estimate (+1 *SD* above the mean of Info Discordance: β = .25, *SE* = .06, *t* = 3.96, *p* < .001). This suggests that people with unfavorable gun attitudes are more likely to update towards the true crime rate when they underestimate crime relative to the social estimate, but more likely to update away from the truth when they overestimate crime relative to the social estimate.

Figure 4.

Gun Attitudes Undermine The Wisdom of Crowds

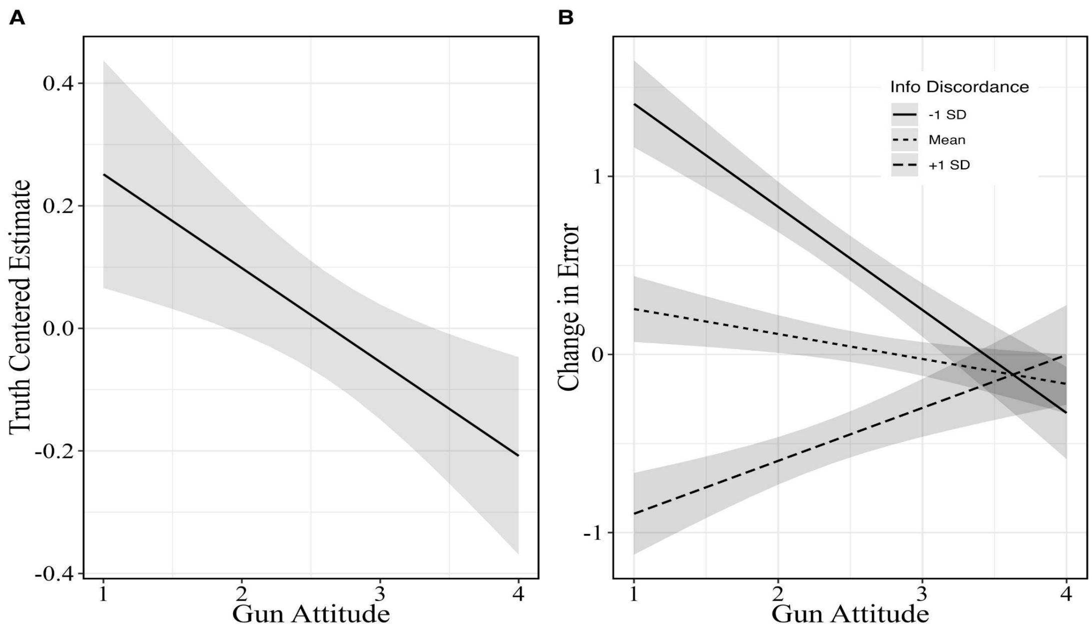

Note. Participants with unfavorable gun attitudes overestimated crime, whereas those with favorable gun attitudes underestimated crime (A). The Influence of social information depends on individual differences in gun attitudes and information discordance (B). Lower gun attitude scores reflect unfavorable gun attitudes whereas higher values reflect favorable gun attitudes. Positive values on the Y-axis of panel A reflect over estimation of crime relative to truth. Positive values on the Y-axis of panel B reflect lower error after viewing social estimates whereas negative reflect more error. Error bars denote 89% CI around the mean.

Discussion

In Study 1, we found that the wisdom of crowds was enhanced after partisans sampled empirical data and furthermore after viewing the average estimate of others compared to when reflecting on their prior answer (control). In addition, we found that individual differences in gun attitudes undermined the wisdom of crowds, as partisans with more extreme gun attitudes deviated further from the truth. Importantly, the social estimates that participants saw in Study 1 were perfect representations of sampled data, thus representing an ideal situation where the wisdom of crowds was unbiased. In contrast, in many real-world contexts, the social estimates that people encounter may be biased by pre-existing beliefs and preferences. Thus, one important question is whether the wisdom of crowds is robust to partisan biased estimates, and whether people achieve greater collective accuracy when they can freely choose the makeup of their samples.

Study 2: Information Propagation & The Wisdom of Crowds

Study 2 aims to explore the progression of the wisdom of crowds over time, examining whether the wisdom of crowds is enhanced or deteriorates when people are free to choose where and how much information they gather. By granting participants the autonomy to determine the composition and makeup of their experiences, we also assess whether the wisdom of crowds is resilient to partisan biases in sampling and evaluations, or whether it introduces systematic biases that entrench polarization. This is particularly relevant given that policy communication frequently spreads across homogenous social networks, where partisans may selectively gather and interpret scientific findings that confirm their pre-existing beliefs, leading to biased interpretations which propagate across social networks and influence others' perceptions and interpretations of the data.

Methods

Participants

Participants (N = 606) were recruited from CloudResearch. The sample size was based on a post-hoc sensitivity analysis from Study 1. All participants self-identified their political identity on a 7-point scale from 1 (very liberal) to 7 (very conservative), with 4 being neither. The task was only made visible to participants who self-identified as either a Democrat or Republican. However, 17 participants still responded with a 4 (i.e., they do not identify as either), and they were excluded from all analyses in line with our pre-registered exclusion criteria. Participants were also excluded if they failed the same manipulation check used in Study 1 (N = 43), as well as if they did not sample at least one time (N = 21). Regarding outliers, participants were excluded if their estimates exceed the true mean by a factor of 1000 (N = 12). Upon inspection of visual diagnostics, we observed an additional 22 outliers that were three standard deviations below the mean at Round 2, which we also excluded from all analysis. Note, however, that results examining changes in accuracy between rounds and conditions did not significantly differ with these observations included (see Table S1 in Supplemental Materials). After applying these exclusion criteria, N = 501 participants remained. The sample had a M_{age} of 39.24 (SD = 10.70) (see Table 1 for demographic summary statistics) and was approved by the Institutional Review Board at a large Southern California University.

Procedure

Study 2 employed a similar design to Study 1 but this time participants were randomly assigned to either sample first-hand empirical evidence about the impact of gun access policies on crime rates from government officials (Empirical condition), or sample second-hand social estimates from Democrats and Republicans who had previously completed the task (Social condition, see Figure 5). The cover story and sampling task for the Social condition were identical to Study 1 and the cover stories were matched as closely as possible (for full cover story, see Supplemental Materials).

In the Social condition, participants were told that there is an ongoing debate in the U.S. about whether expanding access to guns increases violent crime by increasing the number of people carrying weapons or decreases crime by making it easier for law-abiding citizens to defend themselves from violent criminals. They were informed that 1,000 Democrat and Republican mTurk workers had previously sampled crime rate data from counties that expanded access to guns 1 year prior, and that we collected their crime rate estimates based on this data. Participants were once again provided the benchmark for crime rate prior to expanding access to guns (i.e., 441 to 451 per 100,000 people).

After viewing the cover story but before beginning the sampling portion of the task, participants provided their prior estimate of the new crime rate, along with a confidence rating for their estimate (round 1 estimate). Prior estimates were collected via a free text entry and confidence ratings were collected via a 100-point slider scale, ranging from 0 (no confidence at all) to 100 (total confidence). We also collected categorical predictions indicating whether participants believed that crime would increase, decrease, or stay the same after the policy had been implemented.

On each trial in the Social condition, participants selected to sample a crime rate estimate from a Democrat or Republican mTurk worker, which was represented by a numerical score. They were subsequently given the option to continue sampling, or stop sampling when they felt they could confidently estimate the average crime rate. The Democrat and Republican crime rate estimates were each 100 integers pulled from a normal distribution based on the mean crime rate estimate that Democrats and Republicans provided in Study 1. Specifically, we generated two normal distributions using the average crime rate estimate from Study 1 Democrats (453) and Republicans (437) that participants provided after sampling the empirical data. Regarding the variance of the distributions[5], we chose the largest variance that would maintain a Cohen's d of .5 between the two distributions, which previous work has demonstrated to be noticeably different to people (see Bergh & Lindskog, 2019). Upon sampling, a JavaScript function in Qualtrics retrieved random numbers (with replacement) from the corresponding dataset. In the social condition, the information presented after each sample included the political affiliation of the mTurk worker (Democrat/Republican Supporter) and their anonymous ID, along with that mTurk workers crime rate estimate. After viewing an example trial, participants could freely gather as much information as they liked until they felt confident to estimate the new crime rate.

[5] See Table 2 in-text for actual mean and standard deviation of estimates as a function of political affiliation.

Once participants decided to stop sampling, they provided a new estimate of the crime rate via a free entry, along with a confidence rating for their estimate (round 2 estimate), in addition to stating whether they thought crime had increased, decreased, or stayed the same.

Critically, all participants were financially incentivized to provide accurate estimates of the crime rate. On top of the payment provided for completing the study, participants were also told that they were eligible to win bonus money based on the accuracy of their final estimates. Ultimately, all participants were awarded the same $.50 bonus at the end of the task, regardless of task performance. Participants were fully debriefed at the end of the study and were informed that the crime rates they saw were not real and were made up for the purpose of the study.

Figure 5.

Task Diagram for Study 2

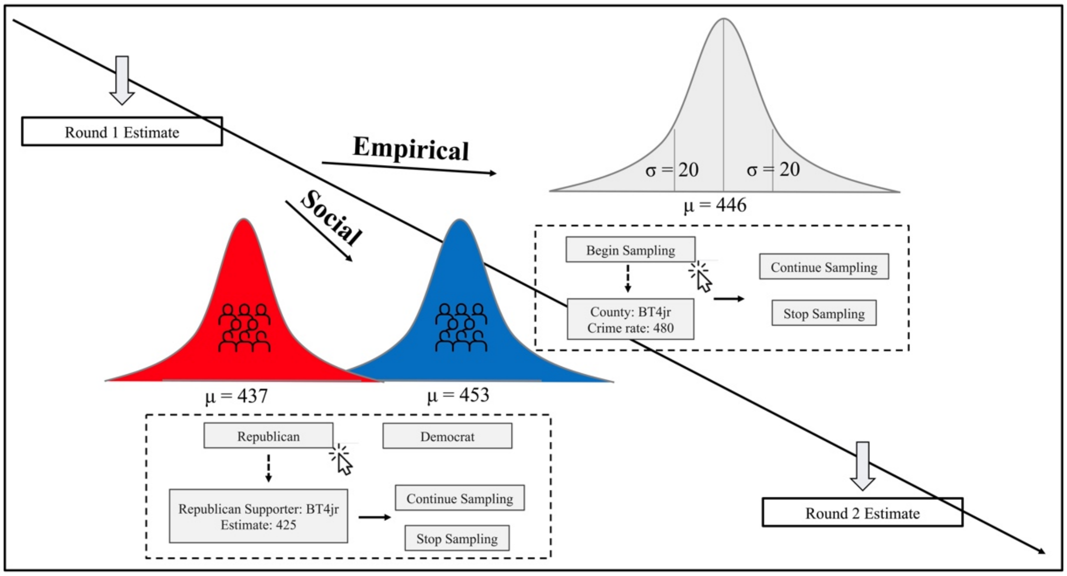

Note. Participants were provided with a baseline for crime rate for the five years prior to the policy change. They were then randomly assigned to freely gather first-hand (empirical) estimates or second-hand (social) estimates from other Democrat and Republicans. There was a total of two rounds of estimates in Study 2.

Self-report measures

Political Identification. Identification with an ideological label was operationalized using a standard 7-point measure of ideological self-description: 1 (*very conservative*), 2 (*conservative*), 3 (*slightly conservative*), 4 (*moderate*), 5 (*slightly liberal*), 6 (*liberal),* and 7 (*very liberal*).

Gun Attitudes Scale (GAS). To assess gun attitudes, participants completed the same 9-item general attitude towards gun scale as participants in Study 1 (Tenhundfeld, 2017). The scale again demonstrated excellent reliability in the current sample ($a = .95$; Omega = .97).

Analysis Plan

Measuring Ingroup & Outgroup Sampling Dynamics. We pre-registered two models interrogating sampling biases, relating to where people sampled first and whether people sampled more from one group over the other. Regarding the first sample, we conducted a binomial test to examine whether the proportion of participants who pick the ingroup category first deviates from chance selection. Regarding overall sampling, we based the decision of whether to use a negative binomial or Poisson generalized linear mixed-effect model based on descriptive statistics of the mean and variance of the data as well as measures of overdispersion using a likelihood ratio test. Results indicated that there was no overdispersion in the current sample (see Table S5 in Supplemental Material), and therefore we report a Poisson generalized linear mixed model. The model included two fixed effects, one for group category (Ingroup, Outgroup), and one for party membership (Democrat vs. Republican) as well as a random effect for subjects.

Estimating Differences in Sampling Behavior Between Conditions and Political Affiliation. To test whether participants sample more in the Empirical or Social condition, as well as whether there were differences in sampling behavior as a function of political affiliation, we report an exploratory negative binomial regression model. This model regresses the total number of samples participants gathered onto a categorical factor representing condition (with the Empirical condition coded as the reference group) and a dummy coded factor representing political affiliation (with Democrat coded as the reference group).

Association Between Sampling Behavior and the Variance of Experiences. To test whether sampling behavior generates greater ingroup relative to outgroup variability, we pre-registered a linear regression comparing the standard deviation of people's ingroup relative to outgroup experiences.

Association Between Sampling Biases and The Wisdom of Crowds. Given that participants in the Social condition had the freedom to choose where to gather information, we conducted several exploratory analyses to determine whether the wisdom of crowds was resilient to biases in sampling behavior.

We first tested whether sampling from the ingroup first was associated with more biased evaluations. To that end, we conducted a Wilcoxon rank-sum test with continuity correction comparing absolute deviations from truth for participant's round 2 estimates as a function of first sample choice. Next, we tested whether participants who sampled overall more information from their own group likewise ended up with more biased estimates, and whether this was moderated by political affiliation. For example, if participants sample more information from ingroup members, then veridical representations of the sampled information will give rise to systematic biases in estimates as Democrat and Republican estimates of the empirical data were biased along party lines (i.e., Democrats overestimate crime and Republicans overestimate crime, on average). To test this, we regressed the log(absolute difference) between participants' estimates and the true crime rate onto a continuous variable representing the ratio of ingroup to outgroup samples, as well as a moderator for political affiliation.

Examining the Influence of First-Hand Versus Second-Hand Information on the Wisdom of Crowds. To examine whether the wisdom of crowds is enhanced following sampling of empirical and social information, we pre-registered two models. The first captures changes in absolute error between participant's round 1 estimate and the true crime rate and participants round 2 estimate and the true crime rate. Specifically, we planned a mixed model that regresses the absolute difference measure onto a categorical variable representing round, with participants coded as a random factor. Upon examination of the model diagnostics, we found that the assumptions of homoscedasticity and normality of residuals were violated (see Figure S2 in Supplemental Materials). To address these issues, we applied a logarithmic transformation to the outcome variable (i.e., absolute difference in error) and included non-logged transformed models in Supplemental materials (see Table S2).

The second model focuses on the proportion of people that update towards the true crime rate which we calculated by creating a variable where 1 reflects a lower error at round 2 compared to round 1, and 0 reflects no reduction in error. We report both Chi-

Square tests of independence, as well as a logistic regression that regresses the binary variables representing learning onto a categorical factor for condition. These analyses determine whether, on average, more people update towards the true crime rate in the Empirical compared to Social condition.

Estimating the Influence of Gun Attitudes on the wisdom of crowds. To examine whether gun attitudes undermine the wisdom of crowds, we pre-registered a linear regression model regressing the difference between participant's round 2 estimates and the true crime rate onto a continuous measure of gun attitudes.

Examining Partisan Bias in Ingroup and Outgroup Estimates. To examine whether partisans were biased in their estimates of ingroup and outgroup averages, we pre-registered a linear regression model that regresses the absolute difference between participant's ingroup and outgroup estimates and the true ingroup and outgroup average, with participants coded as a random factor.

To capture the directionality of the effect, we pre-registered separate linear regression models examining deviations from the true Democrat average onto a measure of political affiliation as well as deviations from the true Republican average onto a measure of political affiliation. To account for non-independence in participants' estimates of the ingroup and outgroup (participants provide estimates for both the ingroup and outgroup), we report a single mixed model regressing the difference between participant's group estimates and the true group averages onto a dummy coded factor representing estimated group (with the ingroup coded as the reference group) and a dummy coded factor representing political affiliation (with Democrat coded as the reference group).

Results

Sampling Behavior

People Sample First and Most Often from Their own Group. We first examined whether most participants sampled from the ingroup first. In line with our prediction, the majority (75%) chose to sample the first piece of information from their own group (binomial test H0 = .5: $p < .0001$, 95% *CI* [69%, 80%], see Figure 6A).

We also predicted that participants would gather overall more information from their own group. Results from a negative binomial mixed model revealed that participants sampled overall more information from their ingroup (*M* = 6.57, *SD* = 6.37, *Mdn* = 5) compared to outgroup (*M* = 5.95, *SD* = 5.09, *Mdn* = 4) ($\beta = -.13$, 95 % *CI* = [-.21 -.06], *SE* = .03, $z = -3.86$, $p < .001$, $sr^2 = .006$, see Figure 6B)[6]. These findings are consistent with past work demonstrating that people tend to be biased in their sampling behavior, sampling first and most often from ingroup members (Bergh & Lindskog, 2019; Derreumaux et al., 2022).

[6] We failed to find any evidence for differences in sampling as a function of political group membership, as tested by including an interaction term between sampled group and a dummy coded factor representing affiliation ($\beta = .002$, *SE* = .07, $z = .036$, $p = .97$).

Figure 6.

Sampling as a Function of Group & Condition

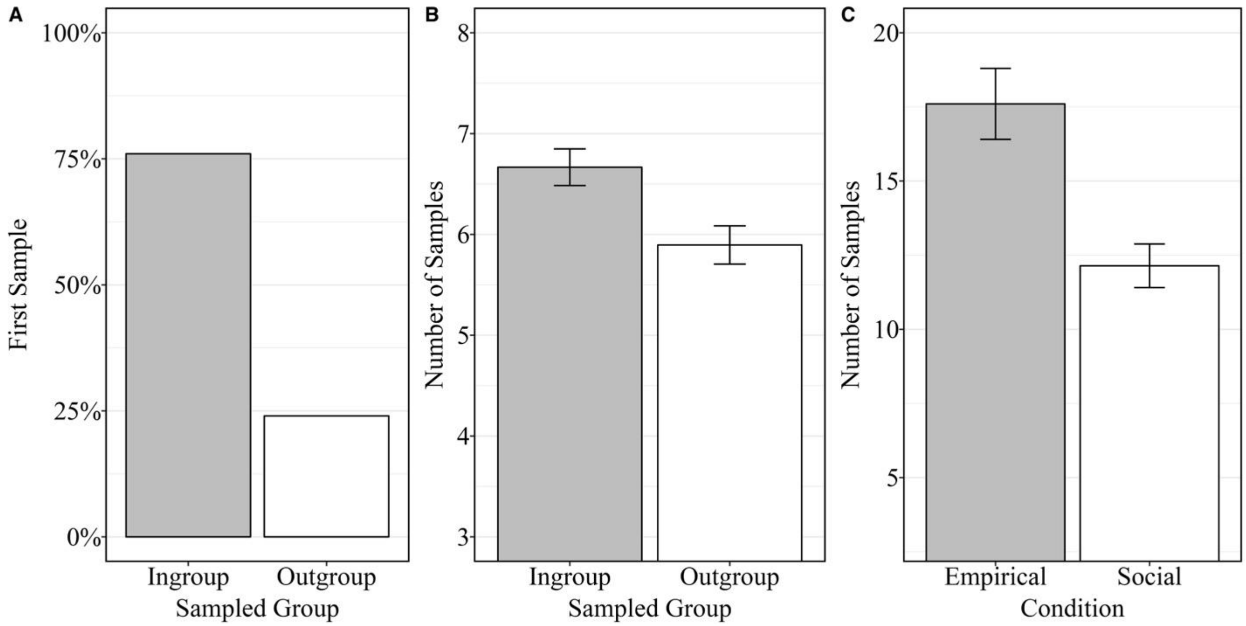

Note. Panel A reflects the proportion of initial ingroup and outgroup samples. Panel B reflects the average sampling from the ingroup and outgroup. Panel C reflects the average sampling in the empirical and social condition. Error bars denote standard error of the mean.

People Sample More in the Empirical Compared to Social Condition. As an exploratory analysis, we also examined whether people sample more information when sampling empirical data compared to sampling social estimates, and whether this depended on participants' political affiliation. We observed a significant main effect of condition, such that participants sampled more in the Empirical compared (M = 17.36, SD = 17.97, Mdn = 13) to Social condition (M = 12.07, SD = 10.97, Mdn = 10) (β = -.43, 95 % CI = [-.62 -.24], SE = .09, z = -4.38, p < .0001, see Figure 6C). In addition, we observed a significant main effect of political affiliation, such that Democrats (M = 16.65, *SD* = 17.65, *Mdn* = 12) sampled overall more than Republicans (M = 12.63, SD = 11.53, Mdn = 9) (β = -.32, 95 % CI = [-.51 -.12], SE = .09, z = -3.27, p = .001). We did not observe a condition by affiliation interaction (β = .15, SE = .13, z = 1.13, p = .25).

Connecting Sampling Behavior to Downstream Estimates

Increased Samples Improves The Wisdom of Crowds. We also conducted an exploratory analysis to determine whether those who sampled more information had more accurate final estimates on average. To that end, we regressed log deviations from the true crime rate onto a variable representing the total number of trials each participant sampled collapsed across conditions. This model revealed a significant main effect of number of trials on accuracy (b = -.01, 95 % CI = [-.02 -.004], SE = .003, t = -3.101, p = .002), demonstrating that people who sample more end up with estimates that are closer to the truth. We did not observe a significant condition by number of trials interaction (b = -.0006, SE = .008, t = -.07, p = .94).

No Evidence that Sampling Behavior Influences the Variability of Experiences. We predicted that biased sampling behavior in the Social condition (i.e., sampling first and most often from the ingroup) would give rise to more variable ingroup experiences. However, we found no evidence that sampling behavior impacts the variance of people's experience in the current sample (β = .02, SE = .08, t = .28, p = .77).

One potential explanation for why we failed to replicate this effect in the current sample is because in our past work, first samples were manipulated to be overly positive or negative. Because most participants sampled from the ingroup first, this introduced systematic variance into people's ingroup experiences. However, we might expect ingroup experiences to be more variable regardless of whether first samples were more variable or not, as past work finds that lower samples underestimate variance and thus more sampling from the ingroup should give rise to more variable ingroup relative to outgroup experiences (i.e., “ingroup heterogeneity effect:” Konovalova & Le Mens, 2020).

Sampling from the Ingroup First Undermines The Wisdom of Crowds. We first tested whether sampling from the ingroup first was associated with more biased evaluations. A Wilcoxon rank-sum test revealed that people who sampled from the ingroup first ended up with estimates that deviated more from the true crime rate (absolute $M = 39.03$, $SD = 77.92$) compared to people who sampled from the outgroup first (absolute $M = 27.69$, $SD = 62.30$) ($W = 7470.5$, $p = .042$).

Sampling More Information from the Ingroup Undermines The Wisdom of Crowds. Next, we examined whether participants who sampled overall more information from their own group ended up with more biased estimates, and whether this was moderated by political affiliation. Results from a linear regression revealed a weak yet significant sampling ratio by affiliation interaction ($b = .59$, 95 % CI = [.01 1.17], SE = .29, $t = 2.02$, $p = .043$), suggesting that as the ratio of ingroup to outgroup samples increases, Republican’s estimates deviated more from truth compared to Democrats. These results indicate that there may be differences in the impact of sampling biases on Democrats and Republicans. For example, Republicans sample overall less than Democrats and therefore sampling biases may have a larger influence on Republican experiences compared to Democrats.

Evaluations

The Influence of First-Hand Versus Second-Hand Information on The Wisdom of Crowds. We first tested whether sampling any information enhances the wisdom of crowds. To that end, we regressed the absolute difference between participants' estimates (log transformed) and the true crime rate onto a variable representing round. Indeed, we found that collective judgements were more accurate after participants sampled information, representing a roughly 83% reduction in absolute error (b = -1.63, 95 % CI = [-1.77 -1.48], SE = .07, t = -20.82, p < .0001, $sr2$ = .24, see Figure 7).

Next, we examined the impact of first-hand versus second-hand knowledge on the wisdom of partisan crowds. To that end, we included an interaction term between round and condition, which revealed a significant round by condition interaction, (b = .43, 95 % CI = [.12 .74], SE = .15, t = 2.78, p = .005, $sr2$ = .006, see Figure 6)[7], demonstrating that collective error was reduced more after partisans sampled first-hand empirical information (87.6% reduction in error) compared to when partisans sampled second-hand social estimates (75% reduction in error). This suggests that while learning from both first-hand and second-hand information enhances the wisdom of crowds, the effect is greater when people sample empirical data directly.

We also examined whether more people updated their estimates towards the true crime rate in the Empirical compared to Social condition. A Chi-Square test of

[7]We included political affiliation as a moderator and failed to find evidence that differences in error between rounds and conditions depended on affiliation (β = .13, SE = .31, t = .43, p = .665).

independence revealed a significant effect of condition on updating, $\chi^2(1) = 5.61$, $p = .017$, demonstrating that more people updated in the correct direction in the Empirical compared to Social condition. Taken together, these results provide converging evidence that the wisdom of crowds is enhanced when people can freely gather first-hand information compared to second-hand social estimates.

Figure 7.

Influence of Information Source on Deviations from Truth

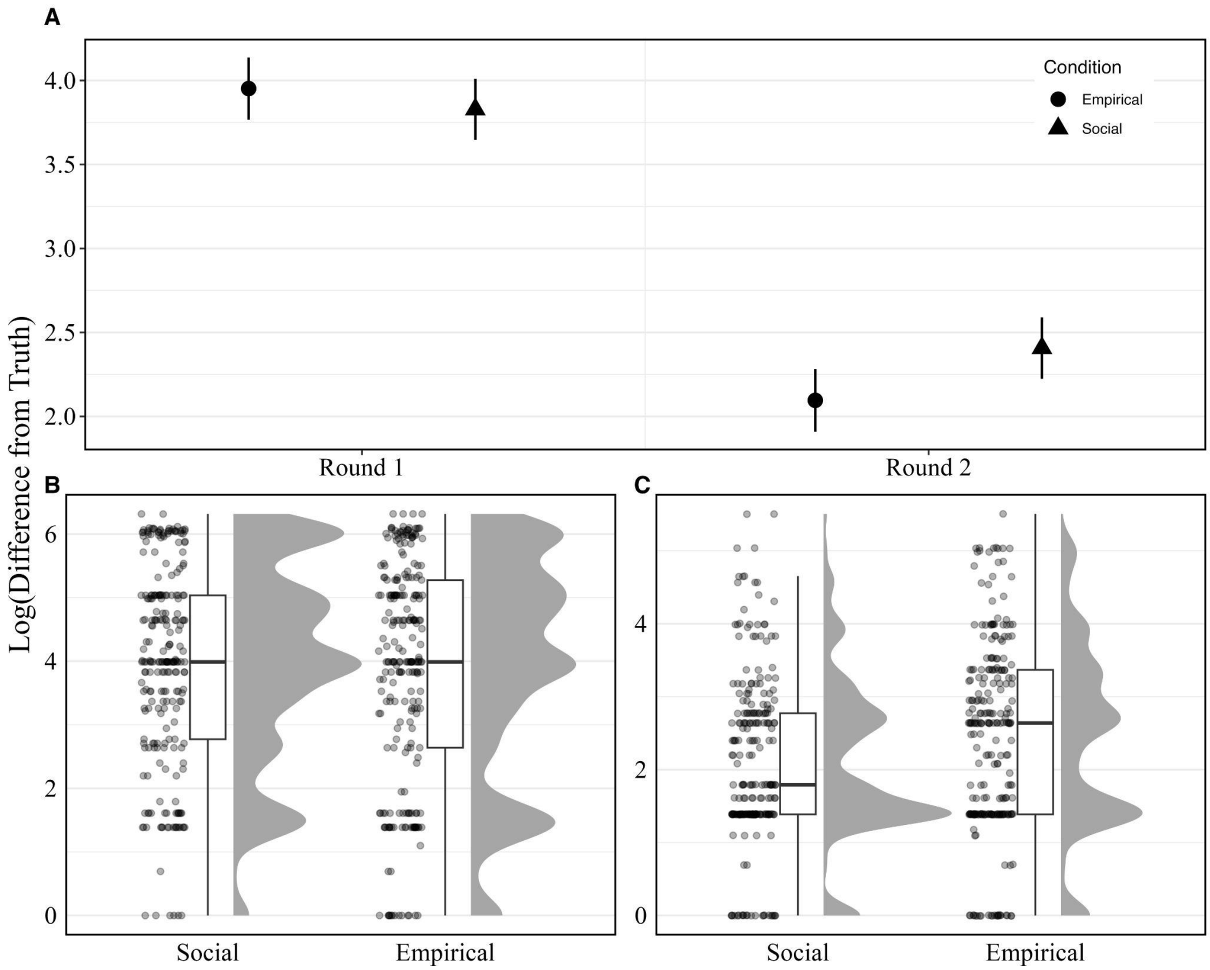

Note. Panel A reflects predicted estimates from a regression model examining log(absolute difference from truth) across rounds and conditions. Panel B reflects raincloud plots showing the distribution of estimates at round 1. Panel C reflects raincloud plots showing the distribution of estimates at round 2. Error bars in Panel A denote 89% confidence intervals.

Extreme Gun Attitudes Undermine The Wisdom of Crowds. We next sought to replicate Study 1 and examine whether people are biased in their interpretation of the sampled information, providing estimates that align with their pre-existing beliefs at the cost of collective accuracy, and furthermore whether this depended on whether partisans sampled first-hand or second-hand information.

Replicating Study 1, we observed a significant negative effect of gun attitudes on deviations from truth (β = -.18, 95 % *CI* = [-.26 -.09], *SE* = .04, *t* = -4.11, *p* < .0001, see Figure 8A)[8], suggesting that partisans with unfavorable gun attitudes overestimated the true crime rate, while those with favorable gun attitudes underestimated it.

Next, we conducted an exploratory analysis to determine whether the influence of gun attitudes on learning differed between the Empirical and Social condition. To that end, we regressed the difference between participants' estimates and the ground truth between rounds (i.e., before and after participants sampled) onto a measure of gun attitudes as well as condition. This model revealed a marginally significant gun attitude by condition interaction (β = -.15, 95 % *CI* = [-.32 .02], *SE* = .08, *t* = -1.75, *p* = .08, see Figure 8B). Simple slopes analysis revealed that people in the Social condition with unfavorable gun attitudes overestimated the true crime rate while participants with favorable gun attitudes underestimated it (β = -.26, *SE* = .06, *t* = -4.29, *p* < .001). However, we found no significant association between gun attitudes and deviations from truth for those in the Empirical condition (β = -.10, *SE* = .06, *t* = -1.58, *p* = .11). These results suggest that gun attitudes may bias information processing more when people sample second-hand social estimates relative to first-hand empirical data.

---

[8] We observed a similar effect of political affiliation, demonstrating that Democrats overestimated the true crime rate, whereas republicans underestimated it (β = -.31, 95 % *CI* = [-.48 -.13], *SE* = .08, *t* = -3.54, *p* < .001).

Figure 8.

Gun Attitudes Undermine the Wisdom of Crowds

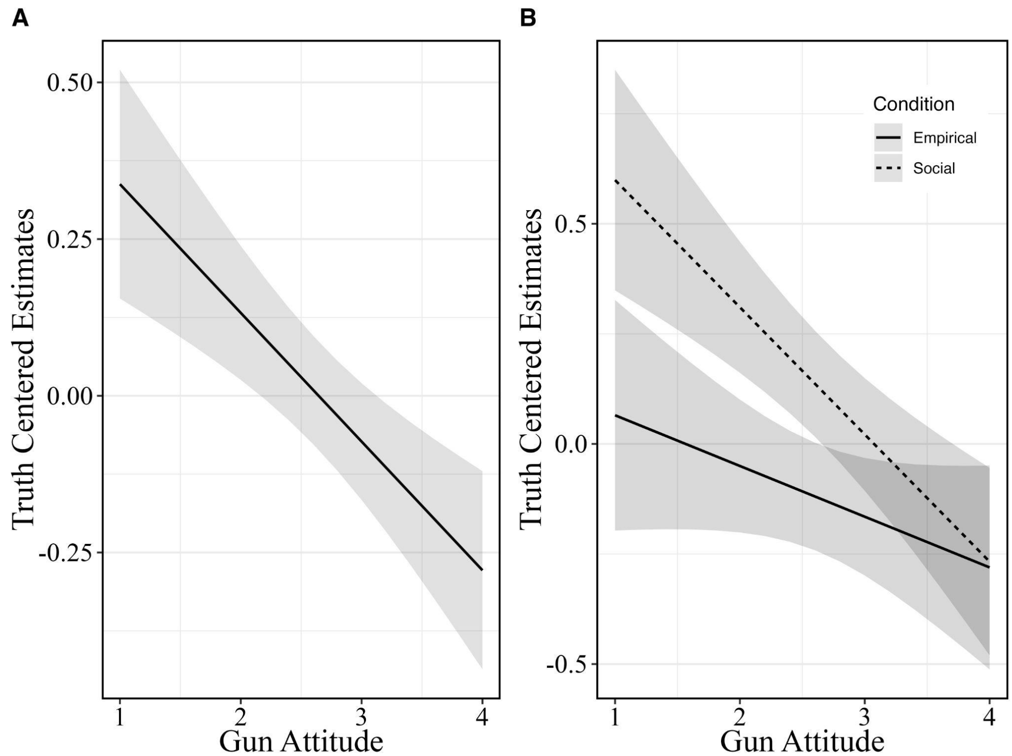

Note. Panel A reflects deviations from truth as a function of gun attitudes. Panel B reflects deviations from truth as a function of gun attitudes and condition. Lower gun attitude scores reflect unfavorable gun attitudes whereas higher values reflect favorable gun attitudes. Positive values on the Y-axis reflect over estimation of crime relative to truth. Error bars denote 89% confidence intervals.

Partisans Exaggerate Group Difference. Next, we examined whether participants were more accurate in their estimates of ingroup compared to outgroup averages. To that end, we regressed the absolute difference between participants' ingroup, and outgroup estimates and the true ingroup and outgroup average onto a dummy coded factor representing the estimated group (i.e., ingroup or outgroup). In contrast to our pre-registered hypothesis, we observed no difference in error as a function of whether people were estimating the ingroup or outgroup (β = .02, *SE* = .05, *t* = .47, *p* = .63). In other words, participants were not more or less accurate in their estimates of the average ingroup and outgroup estimate.

Next, we sought to test the direction of partisans' ingroup and outgroup estimates and whether partisans were more likely to over or under-estimate the true group average based on their affiliation. This model revealed an estimated group by affiliation interaction (β = .30, 95 % *CI* = [.06 .55], *SE* = .12, *t* = 2.48, *p* = .01, see Figure 9). Post hoc simple contrasts revealed that Democrats tended to overestimate the true Democrat average and underestimate the true Republican average (*t*(115) = 1.74, *d* = .16, *p* = .08), whereas Republicans tended to underestimate the true Republican average and overestimate the true Democrat average (*t*(136) = -1.74, *d* = -.13, *p* = .08). This suggests that partisans may exaggerate the differences they perceive between both ingroup and outgroup members in line with general stereotypes about Democrats and Republicans, in this case, that Democrats will be against gun access and overestimate crime after the policy change and Republicans will be in favor of gun access and underestimate crime after the policy change.

Figure 9.

Partisans Exaggerate Group Differences

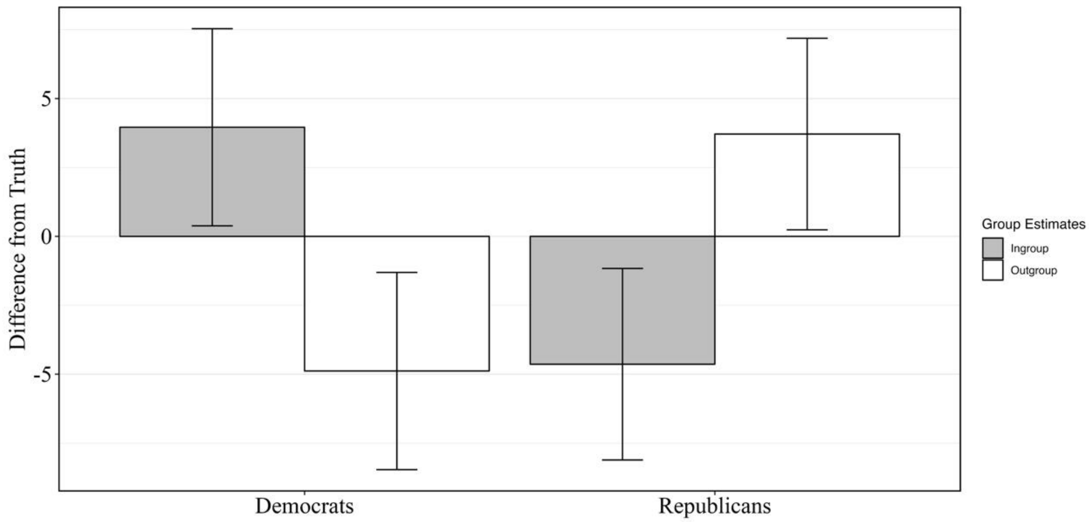

Note. This figure shows the average difference between ingroup and outgroup estimates and the true ingroup and outgroup average as a function of political affiliation. Error bars denote standard error of the mean.

Discussion

Study 2 demonstrates the wisdom of crowds is enhanced when partisans have access to first-hand (Empirical) and second-hand (Social) information. Importantly, while both contexts enhanced the wisdom of crowds, the effect was attenuated in the Social condition, suggesting that social information processing undermines the wisdom of crowds relative to contexts where people have direct access to empirical data.

A potential explanation for the observed learning decrements in the Social condition may be linked to participants having the freedom to choose the composition of their networks, which determined where and how much information they gathered. Given that Democrats typically overestimate crime while Republicans underestimate it, individuals who produce veridical estimations of sampled information (i.e., estimates that accurately reflect their experiences) will display biased judgments if their samples are not representative. In fact, we found that people exhibited biased sampling behavior when sampling social estimates, as they tended to sample first and most often from their in-group, which in turn was associated with more biased evaluations. Collectively, these results suggest that Democrats and Republicans possess a remarkable capacity to learn from their hands-on experiences; however, the wisdom of crowds is compromised when individuals lack access to first-hand information and consequently gather unrepresentative samples from their communities.

Replicating Study 1, we also observed that gun attitudes impede the wisdom of crowds, with individuals holding unfavorable gun attitudes overestimating the true crime rate, while those with favorable gun attitudes underestimate it. Notably, the impact of gun attitudes on the wisdom of crowds was more pronounced for participants who sampled second-hand information compared to those who sampled first-hand information. This suggests that pre-existing beliefs about guns may contribute to a more biased integration of social information. However, when individuals have access to first-hand empirical data, they are less likely to exploit the sampled information and arrive at biased conclusions, suggesting a potential avenue for mitigating the influence of pre-existing beliefs on learning outcomes.

General Discussion

The primary goal of this dissertation was to examine the interplay between the ways partisans gather information and update their beliefs about contentious issues, and to examine the implications of this process on the wisdom of crowds. Additionally, this research sought to determine whether providing individuals with a diverse range of experiences that draw from both first-hand experiences and second-hand social estimates from one's communities could mitigate polarization and enhance the wisdom of partisan crowds. To that end, we integrated research on sampling models with motivated reasoning to investigate the various conditions under which the wisdom of crowds persists, as well as the extent to which individuals leverage or exploit social knowledge to arrive at favorable conclusions.

Throughout our studies, partisans learned about the impact of a policy that increased access to guns on subsequent crime rates. We manipulated their access to first-hand empirical information or second-hand social estimates. In contrast to previous work examining the impact of social influence on the wisdom of crowds, where partisans are exposed to the average estimates of a small number of partisans (see e.g., Becker et al., 2017, 2019; Guilbeault et al., 2018; Jayles et al., 2017; Lorenz et al., 2011; Mavrodiev & Schweitzer, 2021), our study design allowed participants to freely choose the makeup of their communities. Furthermore, participants were provided with accurate prior knowledge in the form of baseline crime rate statistics. By enabling participants to decide where and how much information to gather, our approach emulates the way people sample and integrate information in their daily lives, including data processed and interpreted by other members within their communities.

Our findings demonstrate that aggregate estimates become more accurate when partisans have access to empirical data (Study 1), and that exposing people to the average estimates of others further enhanced the wisdom of crowds (Study 1). Despite reducing collective error on average, crime rate estimates were systematically biased towards people’s pre-existing gun attitudes, such that those with unfavorable gun attitudes overestimated the true crime rate, while those with favorable gun attitudes underestimated it. Extending past work demonstrating that the wisdom of crowds is robust to partisan social influence in decentralized networks (Becker et al., 2019), we find that mean responses become more accurate as a result of social influence, even in contexts where people can freely choose the makeup of their communities. However, while social influence did improve mean responses, they were not as accurate as when people had access to first-hand information (Study 2). This was attributable to (a) biases in people's sampling behavior of social information, as partisans sampled from their in-group first and most often, and (b) given that ingroup and outgroup estimates were systematically biased, participants experiences were biased to the extent that they were uneven in their sampling behavior.

Taken together, our research highlights the importance of providing partisans with first-hand empirical knowledge to enhance collective judgements and reduce polarization. Moreover, we illuminate the boundary conditions of partisan social influence in decentralized networks on the wisdom of crowds, demonstrating that although social influence can improve collective judgements to a certain extent, it may also lead to systematic partisan biases that further entrench partisan divides, particularly when the social information that partisans sample propagates far from the empirical data (e.g., third and fourth-generation estimation).

Informed Decision Making Enhances The Wisdom of Crowds

Research examining the impact of partisan social influence on collective judgements has primarily relied on contexts where people are asked to provide independent estimates on niche statistics that they know nothing about. Due to these uninformed priors, initial estimates tend to include a great deal of error and vary widely across samples. Thus, even when individual estimates are systematically biased (e.g., due to partisan biased interpretations), averaging estimates brings people closer to the truth, on average, in line with the wisdom of crowds. In the current work, we provide partisans with contextual information about the data before they give estimates in the form of the range of crime rates one year prior to the policy change. This baseline provides a benchmark with which to compare new crime rates. We find that allowing partisans to freely sample first-hand (empirical) data about crime rates dramatically increased the wisdom of crowds (Study 1), even in contexts where people have informed priors. Moreover, when given the opportunity (e.g., sample individual or view average estimates of others), partisans tapped into the wisdom of crowds, further enhancing the accuracy of their own judgements (Study 2). Notably, however, we also observed systematic biases in crime rate estimates based on people’s pre-existing attitudes and beliefs, such that those with unfavorable gun attitudes over estimating the true crime rate and those with favorable gun attitudes underestimating it.

The current findings extend past research by examining how social influence functions in everyday settings where individuals have varying levels of prior knowledge and are exposed to the estimates of other partisans within their social networks. By incorporating informed priors into our research design, we offer a more realistic representation of the complexities faced in daily life, and contribute valuable insights into the dynamics of the wisdom of crowds and its potential for mitigating biases in decision-making processes. These results are encouraging as they suggest the wisdom of crowds is robust to contexts where people have informed priors, yet it is important to consider that all partisans were provided with the same accurate information regarding the prior range of crime statistics. As such, our findings speak to an ideal scenario where partisans are provided with a common source of knowledge that closely approximates reality. However, in many real-world situations, people may not be informed at all, or rely purely on their associations or stereotypes, or worse, they may be provided with inaccurate disinformation designed to mislead. Future research should account for idiosyncrasies in people's prior knowledge, and examine the impact of social influence on the wisdom of crowds in contexts where people are free to select the source of their prior knowledge.

Social Influence and Information Propagation

The current research extends theory on the wisdom of crowds by examining how individuals learn from others within a sampling framework. Previous research manipulated social influence via prescribed decentralized networks of a few partisans (Becker et al., 2017, 2019; Guilbeault et al., 2018), and found social influence improved the accuracy of group estimates, even when participant's beliefs became more similar. We extend this work by allowing partisans to freely choose where and how much information to gather. By granting participants the autonomy to determine the composition and makeup of their communities, we assess the resilience of the wisdom of crowds against biases in individuals' sampling and interpretation of social information. Furthermore, by comparing how partisans' sample and evaluate first-hand versus second-hand information, we shed light on how disagreements may arise as a function of the information sources that people gather from.

We found that people sampled overall more information when gathering first-hand compared to second-hand information (Study 2), and that the wisdom of crowds was enhanced when people sampled first relative to second-hand information. Theoretically, people may require fewer second-hand compared to first-hand samples, as each sampled estimate itself comprises multiple samples, akin to a sampling distribution. However, we found that people were biased in their sampling behavior, sampling first and most often from their ingroup, leading to unrepresentative samples. Moreover, people who sampled from the ingroup first ended up with the most biased estimates, but people who sampled from the outgroup first had final estimates that were no less accurate than people who sampled first-hand empirical information. These findings suggest that biases in sampling of social information may underlie learning decrements, and that changing the social information people seek out may provide one route to increasing collective accuracy. In ongoing research, we build on this finding by allowing people to sample both first-hand and second-hand information, but manipulate the order in which information is sampled. One possibility is that providing people with first-hand empirical information first can inoculate against biases that arise when people sample second-hand information. Alternatively, providing people access to first-hand knowledge at any stage during information processing may serve as a protective measure against partisan-biased social estimates.

Interestingly, while social learning did improve collective accuracy on average, our results suggest that it may also introduce systematic partisan biases that further entrench partisan divides, particularly when social estimates propagate further from the empirical data (e.g., third and fourth-generation estimation). For instance, we found that when partisans gathered first-hand (empirical) information, they interpreted the data as being consistent with their prior beliefs and attitudes, with Democrats overestimating and Republicans underestimating the true crime rate. Subsequently, partisans who gathered these second-hand estimates further exaggerated group differences, perceiving Democrats as more extreme over-estimators and Republicans as more extreme under-estimators of crime than they were. This suggests that although social learning may reduce error to a certain extent, achieving accurate collective judgments based purely on social information may be challenging to the extent that people are biased in their sampling behavior, even in decentralized networks.

In future research, we will conduct a third wave of sampling to test this empirical question. Specifically, participants will be randomly assigned to either gather first-hand empirical information or third-hand social estimates collected in Study 2. One possibility is that as social information propagates farther from the empirical data, more socially digested beliefs may undermine the wisdom of crowds and increase polarization. In contrast, if partisans exhibit even-handed sampling behavior (e.g., sampling evenly from Democrats and Republicans) and provide estimates that accurately represent the sampled information, then proximity to the empirical data may not be essential for the wisdom of crowds. Examining whether collective judgments deviate farther from the truth following social information processing of third-hand information will provide insight into the boundary conditions of social influence in the wisdom of crowds. By understanding these dynamics, we can better inform strategies to mitigate the negative effects of polarization and promote more accurate collective decision-making processes in various contexts.

Moving Beyond Collective Judgements of Crime Rate Statistics

The gun control debate stands as one of the most polarizing issues in recent U.S. history, with political gridlock in Washington reflecting unwavering support among constituents, offering limited opportunities for dialogue and compromise. This debate is further complicated by challenges in estimating the causal effects of gun access policies on crime rates due to policy spillover and gun migration (Coates & Pearson-Merkowitz, 2017). Due to this impasse, it is encouraging to observe that the wisdom of crowds prevails when partisans learn about crime rate statistics, especially when they have access to first-hand empirical data.

To build upon our findings, future research could investigate the influence of prior knowledge and sampling biases on the wisdom of crowds in different polarizing contexts that vary in issue features, such as climate change and immigration. Perceptions of threat surrounding a given issue may differentially influence how Democrats and Republicans respond to social influence, potentially influencing how people learn from first-hand versus second-hand information, leading to greater bias in one group over another. For example, researchers could examine the role of perceived threat associated with a given issue in shaping the responsiveness of Democrats and Republicans to social influence. Perceived threat exerts a strong influence over cognition, influencing the information processing strategies people employ (Dawson et al., 2002) and how people learn from ingroup and outgroup members (Derreumaux, Elder, et al., 2023; Golkar & Olsson, 2017). One possibility is that when a particular issue is perceived as a high threat by one group, it may heighten their motivation to adhere to group norms, leading to greater adoption of ingroup estimates and resistance to opposing estimates, thereby undermining the wisdom of crowds. In contrast, if a group perceives an issue as a lower threat, they may be more open to diverse perspectives and social learning. One important consideration is that perceived threat between groups may be heightened in a political context characterized by competition and zero-sum contest over success and failure (Brewer, 1979; Chang et al., 2016; Cikara et al., 2017). A thorough investigation of how these intergroup dynamics play out under different threat conditions and in different groups may help in designing interventions that increase social information processing and bridge intergroup divides.

Another important consideration is examining partisan differences in trust towards different institutions and information sources writ large (see Druckman & McGrath, 2019 for review), and how mistrust in these institutions might impact how people gather and evaluate information. Partisans differ in their trust towards government institutions (e.g., public health institutions), and may therefore also differ with respect to how they sample empirical data, which will impact the wisdom of the crowd. For instance, Republicans are generally distrustful of research on climate change (Druckman

& McGrath, 2019), but may be more trusting of government institutions that report on crime rate statistics. As such, their initial crime rate estimates may be more accurate on average, leading to more even-handed sampling from both Democrats and Republicans and ultimately more accurate collective judgements. In contrast, greater distrust may give rise to biased sampling and evaluations of both empirical and social information, undermining the wisdom of crowds. Understanding how perceptions of trust influence the uptake of information could help in tailoring communication strategies and incentivize social learning to bridge the gap between opposing groups, ensuring that the information shared is perceived as credible and reliable by both sides.

Network Centralization, Financial Incentives & Polarization

Our results reveal that information exchange and social influence can bolster the wisdom of crowds, which may appear to contradict the polarization of public opinion on gun control policies. Recent theoretical advancements highlight that the structure of networks in which information is communicated is a critical factor in determining the impact of social influence on collective judgments. Specifically, the wisdom of crowds benefits from social information processing within a decentralized network (Centola, 2022). Although participants had the autonomy to select their social network composition, the network structure remained decentralized, ensuring each estimate (and each social connection) carried equal influence. In contrast, many online social networks, such as social media platforms, are inherently centralized, where a select few voices hold disproportionate sway (e.g., reaching hundreds of thousands of followers). This raises important questions regarding the role of status indicators or message approval, such as likes and retweets, in shaping how partisans collect and assess second-hand judgments in group-based decision-making contexts. For instance, in online forums or discussion boards featuring diverse perspectives on gun control policies, a centralized structure may result in users with more followers receiving additional upvotes or endorsements, causing their posts to be prioritized over lesser-known users' contributions, which could ultimately undermine the wisdom of crowds.

To address these limitations and questions, future research could explore the impact of different network structures on the wisdom of crowds. For example, researchers could experiment with manipulating the visibility of status indicators (e.g., likes and retweets) or the prominence of high-status users within online discussion forums to examine how these factors influence the exchange of information and the formation of group judgments. Additionally, researchers could analyze how different levels of moderation or algorithms designed to promote diverse perspectives within online discussions affect the wisdom of crowds and polarization of opinions on controversial topics like gun control policies. By better understanding the interplay between sampling behavior and social influence as a function of network structure, future research could inform the design of online platforms and discussion spaces that foster more balanced and informed decision-making processes.

When assessing the generalizability of the current findings, it is crucial to consider that participants were offered financial incentives to enhance accuracy, whereas political attitudes are typically formed without such incentives. The rationale behind financially incentivizing accuracy was to engage participants (e.g., encourage sampling) and deter nonsensical responses. While the financial incentive offers a stringent test of partisan bias by juxtaposing directed motivations with accuracy-driven ones, it may also amplify the effect of social influence on the wisdom of crowds. Importantly, previous research indicates that people can improve their accuracy without financial incentives (Wood & Porter, 2019), implying that these incentives might affect the magnitude but not the direction of the effect. Financial incentives could potentially impact the wisdom of crowds by increasing participants' motivation to scrutinize information more critically or by promoting a greater willingness to revise personal opinions in light of new evidence. However, these incentives could also inadvertently introduce bias, as participants may be more likely to conform to the perceived majority opinion in order to secure financial rewards.

Future research could address these limitations by exploring the role of financial incentives the wisdom of partisan crowds. For instance, researchers could examine the effects of varying incentive structures on collective judgments, such as offering no financial incentives, providing incentives based on individual accuracy, or rewarding group performance. These comparisons could help illuminate how different incentive systems influence participants' motivation to prioritize information acquisition from different sources, revise beliefs based on new evidence, and engage in group decision-making processes. Another avenue for future research could involve examining the impact of financial incentives on the quality and diversity of information exchanged within online discussions or social networks. For example, by manipulating the presence and magnitude of financial rewards, researchers could assess whether incentives promote the sharing of higher-quality information, foster more balanced discussions, or potentially contribute to groupthink and conformity.

Concluding Remarks

In today's U.S. society, disagreement over fundamental facts poses significant challenges to constructive dialogue, effective decision-making, and governance. This pervasive divisiveness not only intensifies social tensions but also hinders progress on critical social issues, such as the development and implementation of common-sense gun laws. Despite these challenges, the current research demonstrates that both Democrats and Republicans possess a remarkable capacity to learn from first-hand experiences and second-hand information shared by fellow partisans, thereby enhancing the wisdom of crowds across the political spectrum. These findings underscore the importance of encouraging individuals to diversify the sources of information and the people they engage with within their social networks. Furthermore, they emphasize the need to empower individuals to become better consumers of first-hand empirical data and to advocate for policy reforms (e.g., promoting science and numeracy literacy) that equip people with the ability to discern quality information sources.

As the challenges of polarization and disagreement on facts persist in contemporary society, it is our hope that this research lays the groundwork for developing solutions that bridge divides and improve the quality of public discourse. By fostering a greater understanding of how individuals gather information and update their beliefs, we can pave the way for more informed decision-making and collaborative problem-solving despite extant partisan divides.

References

Albarracín, D., Gillette, J. C., Earl, A. N., Glasman, L. R., Durantini, M. R., & Ho, M.-H. (2005). A Test of Major Assumptions About Behavior Change: A Comprehensive Look at the Effects of Passive and Active HIV-Prevention Interventions Since the Beginning of the Epidemic. *Psychological Bulletin*, *131*(6), 856–897. https://doi.org/10.1037/0033-2909.131.6.856

Almaatouq, A., Rahimian, M. A., Burton, J. W., & Alhajri, A. (2022). The distribution of initial estimates moderates the effect of social influence on the wisdom of the crowd. *Scientific Reports*, *12*(1), Article 1. https://doi.org/10.1038/s41598-022-20551-7

Anglemyer, A., Horvath, T., & Rutherford, G. (2014). The accessibility of firearms and risk for suicide and homicide victimization among household members: A systematic review and meta-analysis. *Annals of Internal Medicine*, *160*(2), 101–110. https://doi.org/10.7326/M13-1301

Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*(4), 390–412. https://doi.org/10.1016/j.jml.2007.12.005

Barneron, M., Allalouf, A., & Yaniv, I. (2019). Rate it again: Using the wisdom of many to improve performance evaluations. *Journal of Behavioral Decision Making*, *32*, 485–492. https://doi.org/10.1002/bdm.2127

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2014). Fitting Linear Mixed-Effects Models using lme4. *ArXiv:1406.5823 [Stat]*. http://arxiv.org/abs/1406.5823

Becker, J., Brackbill, D., & Centola, D. (2017). Network dynamics of social influence in the wisdom of crowds. *Proceedings of the National Academy of Sciences*, *114*(26). https://doi.org/10.1073/pnas.1615978114

Becker, J., Porter, E., & Centola, D. (2019). The wisdom of partisan crowds. *Proceedings of the National Academy of Sciences*, *116*(22), 10717–10722. https://doi.org/10.1073/pnas.1817195116

Bergh, R., & Lindskog, M. (2019). The group-motivated sampler. *Journal of Experimental Psychology: General*, *148*, 845–862. https://doi.org/10.1037/xge0000601

Braman, D., & Kahan, D. M. (2001). *More Statistics, Less Persuasion: A Cultural Theory of Gun-Risk Perceptions* (SSRN Scholarly Paper No. 286205). https://doi.org/10.2139/ssrn.286205

Brewer, M. B. (1979). In-group bias in the minimal intergroup situation: A cognitive-motivational analysis. *Psychological Bulletin*, *86*, 307–324. https://doi.org/10.1037/0033-2909.86.2.307

Brewer, M. B. (2001). Ingroup identification and intergroup conflict: When does ingroup love become outgroup hate? In *Social identity, intergroup conflict, and conflict reduction* (pp. 17–41). Oxford University Press.

Centola, D. (2022). The network science of collective intelligence. *Trends in Cognitive Sciences*, *26*(11), 923–941. https://doi.org/10.1016/j.tics.2022.08.009

Chang, L. W., Krosch, A. R., & Cikara, M. (2016). Effects of intergroup threat on mind, brain, and behavior. *Current Opinion in Psychology*, *11*, 69–73. https://doi.org/10.1016/j.copsyc.2016.06.004

Cikara, M., Van Bavel, J. J., Ingbretsen, Z. A., & Lau, T. (2017). Decoding "us" and "them": Neural representations of generalized group concepts. *Journal of Experimental Psychology: General*, *146*, 621–631. https://doi.org/10.1037/xge0000287

Coates, M., & Pearson-Merkowitz, S. (2017). Policy Spillover and Gun Migration: The Interstate Dynamics of State Gun Control Policies. *Political Science Faculty Publications*. https://doi.org/10.1111/ssqu.12422

Cohen, G. L. (2003). Party Over Policy: The Dominating Impact of Group Influence on Political Beliefs. *Journal of Personality and Social Psychology*, *85*(5), 808–822. https://doi.org/10.1037/0022-3514.85.5.808

Cook, P., & Ludwig, J. (2006). The Social Costs of Gun Ownership. *Journal of Public Economics*, *90*, 379–391. https://doi.org/10.1016/j.jpubeco.2005.02.003

Cramer, C. E., & Burnett, D. (2012). *Tough Targets: When Criminals Face Armed Resistance from Citizens* (SSRN Scholarly Paper No. 2226957). https://papers.ssrn.com/abstract=2226957

Dawson, E., Gilovich, T., & Regan, D. T. (2002). Motivated Reasoning and Performance on the was on Selection Task. *Personality and Social Psychology Bulletin*, *28*(10), 1379–1387. https://doi.org/10.1177/014616702236869

Denrell, J. (2005). Why Most People Disapprove of Me: Experience Sampling in Impression Formation. *Psychological Review*, *112*(4), 951–978. https://doi.org/10.1037/0033-295X.112.4.951

Denrell, J., & Le Mens, G. (2017). Information sampling, belief synchronization, and collective illusions. *Management Science*, *63*, 528–547. https://doi.org/10.1287/mnsc.2015.2354

Derreumaux, Y., Bergh, R., & Hughes, B. L. (2022). *Partisan-Motivated Sampling: Re-Examining Politically Motivated Reasoning Across the Information Processing Stream*. https://doi.org/10.1037/pspi0000375.

Derreumaux, Y., Elder, J., Suri, G., Ben-Zeev, A., Quimby, T., & Hughes, B. L. (2023). Stereotypes disrupt probabilistic category learning. *Journal of Experimental Psychology: General*, No Pagination Specified-No Pagination Specified. https://doi.org/10.1037/xge0001335

Derreumaux, Y., Lindskog, M., Bergh, R., & Hughes, B. L. (2023). Group-Motivated Sampling: From Skewed Experiences to Biased Evaluations. In K. Fiedler, P. Juslin, & J. Denrell (Eds.), *Sampling in judgment and decision making*. Cambridge University Press.

Derreumaux, Y., Shamsian, K., & Hughes, B. L. (2023). Computational underpinnings of partisan information processing biases and associations with depth of cognitive reasoning. *Cognition*, *230*, 105304. https://doi.org/10.1016/j.cognition.2022.105304

Ditto, P. H., Liu, B. S., Clark, C. J., Wojcik, S. P., Chen, E. E., Grady, R. H., Celniker, J. B., & Zinger, J. F. (2019). At Least Bias Is Bipartisan: A Meta-Analytic Comparison of Partisan Bias in Liberals and Conservatives. *Perspectives on Psychological Science*, 19.

Ditto, P. H., & Lopez, D. F. (1992). Motivated skepticism: Use of differential decision criteria for preferred and nonpreferred conclusions. *Journal of Personality and Social Psychology*, *63*, 568–584. https://doi.org/10.1037/0022-3514.63.4.568

Doherty, B. (2008). *Gun Control on Trial: Inside the Supreme Court Battle Over the Second Amendment*. Cato Institute.

Druckman, J. N., & McGrath, M. C. (2019). The evidence for motivated reasoning in climate change preference formation. *Nature Climate Change*, *9*(2), 111–119. https://doi.org/10.1038/s41558-018-0360-1

Feldman, D. C. (1984). The development and enforcement of group norms. *The Academy of Management Review*, *9*, 47–53. https://doi.org/10.2307/258231

Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. *Psychological Review*, *107*, 659–676. https://doi.org/10.1037/0033-295X.107.4.659

Flavell, J. H., & Ross, L. (1981). *Social Cognitive Development: Frontiers and Possible Futures*. CUP Archive.

Galton, F. (1907). Vox Populi (The Wisdom of Crowds). *Nature*, *75, 450-451.*, 7.

Golkar, A., & Olsson, A. (2017). The interplay of social group biases in social threat learning. *Scientific Reports*, *7*(1), 7685. https://doi.org/10.1038/s41598-017-07522-z

Green, P., & MacLeod, C. J. (2016). SIMR: An R package for power analysis of generalized linear mixed models by simulation. *Methods in Ecology and Evolution*, *7*(4), 493–498. https://doi.org/10.1111/2041-210X.12504

Greenberg, J., & Pyszczynski, T. (1985). Compensatory Self-Inflation: A Response to the Threat to Self-Regard of Public Failure. *Journal of Personality and Social Psychology*, *49(1), 273–280*, 8. https://doi.org/10.1037/0022-3514.49.1.273

Guilbeault, D., Becker, J., & Centola, D. (2018). Social learning and partisan bias in the interpretation of climate trends. *Proceedings of the National Academy of Sciences*, *115*(39), 9714–9719. https://doi.org/10.1073/pnas.1722664115

Hart, W., Albarracín, D., Eagly, A. H., Brechan, I., Lindberg, M. J., & Merrill, L. (2009). Feeling validated versus being correct: A meta-analysis of selective exposure to information. *Psychological Bulletin*, *135*(4), 555–588. https://doi.org/10.1037/a0015701

Hemenway, D. (2011). Risks and Benefits of a Gun in the Home. *American Journal of Lifestyle Medicine*, *5*(6), 502–511. https://doi.org/10.1177/1559827610396294

Hogg, M. A., & Reid, S. A. (2006). Social identity, self-categorization, and the communication of group norms. *Communication Theory*, *16*, 7–30. https://doi.org/10.1111/j.1468-2885.2006.00003.x

Howard, J. W., & Rothbart, M. (1980). Social Categorization and Memory for In-Group and Out-Group Behavior. *Journal of Personality and Social Psychology*, *38(2), 301–310*, 10. https://doi.org/10.1037/0022-3514.38.2.301

Hughes, B. L., Zaki, J., & Ambady, N. (2017). Motivation alters impression formation and related neural systems. *Social Cognitive and Affective Neuroscience*, *12*(1), 49–60. https://doi.org/10.1093/scan/nsw147

Jaeger, B., Edwards, L., Das, K., & Sen, P. (2016). An R 2 statistic for fixed effects in the generalized linear mixed model. *Journal of Applied Statistics*, 1–20. https://doi.org/10.1080/02664763.2016.1193725

Jayles, B., Kim, H., Escobedo, R., Cezera, S., Blanchet, A., Kameda, T., Sire, C., & Theraulaz, G. (2017). How social information can improve estimation accuracy in human groups. *Proceedings of the National Academy of Sciences*, *114*(47), 12620–12625. https://doi.org/10.1073/pnas.1703695114

Juslin, P., Winman, A., & Hansson, P. (2007). The naïve intuitive statistician: A naïve sampling model of intuitive confidence intervals. *Psychological Review*, *114*(3), 678–703. https://doi.org/10.1037/0033-295X.114.3.678

Kahan, D. M. (2013). Ideology, motivated reasoning, and cognitive reflection. *Judgment and Decision Making*, *8*(4), 18. https://doi.org/10.1017/S1930297500005271

Kahan, D. M., Peters, E., Dawson, E. C., & Slovic, P. (2017). Motivated numeracy and enlightened self-government. *Behavioural Public Policy*, *1*(1), 54–86. https://doi.org/10.1017/bpp.2016.2

Kahn, C. (2021, May 24). *53% of Republicans view Trump as true U.S. president - Reuters/Ipsos*. https://www.reuters.com/world/us/53-republicans-view-trump-true-us-president-reutersipsos-2021-05-24/

Kleck, G. (2015). The Impact of Gun Ownership Rates on Crime Rates: A Methodological Review of the Evidence. *Journal of Criminal Justice*, *43*(1), 40–48. https://doi.org/10.1007/BF01064462

Kleck, G., & Patterson, E. B. (1993). The Impact of Gun Control and Gun Ownership Levels on Violence Rates. *Journal of Quantitative Criminology*, *9*(3), 249–287.

Konovalova, E., & Le Mens, G. (2020). An information sampling explanation for the in-group heterogeneity effect. *Psychological Review*, *127*, 47–73. https://doi.org/10.1037/rev0000160

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, *108*(3), 480–498. https://doi.org/10.1037/0033-2909.108.3.480

Kunda, Z., & Nisbett, R. E. (1986). The psychometrics of everyday life. *Cognitive Psychology*, *18*(2), 195–224. https://doi.org/10.1016/0010-0285(86)90012-5

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software*, *82*, 1–26. https://doi.org/10.18637/jss.v082.i13

Lacombe, M. J. (2019). The Political Weaponization of Gun Owners: The National Rifle Association's Cultivation, Dissemination, and Use of a Group Social Identity. *The Journal of Politics*, *81*(4), 1342–1356. https://doi.org/10.1086/704329

Lindskog, M., Winman, A., & Juslin, P. (2013). Naïve point estimation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(3), 782–800. https://doi.org/10.1037/a0029670

Lorenz, J., Rauhut, H., Schweitzer, F., & Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences*, *108*(22), 9020–9025. https://doi.org/10.1073/pnas.1008636108

Mavrodiev, P., & Schweitzer, F. (2021). The ambigous role of social influence on the wisdom of crowds: An analytic approach. *Physica A: Statistical Mechanics and Its Applications*, *567*, 125624. https://doi.org/10.1016/j.physa.2020.125624

McCoy, J., Rahman, T., & Somer, M. (2018). Polarization and the Global Crisis of Democracy: Common Patterns, Dynamics, and Pernicious Consequences for Democratic Polities. *American Behavioral Scientist*, *62*(1), 16–42. https://doi.org/10.1177/0002764218759576

Meiser, T., & Hewstone, M. (2006). Illusory and spurious correlations: Distinct phenomena or joint outcomes of exemplar-based category learning? *European Journal of Social Psychology*, *36*, 315–336. https://doi.org/10.1002/ejsp.304

Morgan, M. G. (2014). Use (and abuse) of expert elicitation in support of decision making for public policy. *Proceedings of the National Academy of Sciences*, *111*(20), 7176–7184. https://doi.org/10.1073/pnas.1319946111

National Center for Health Statistics,. (2021). Centers for Disease Control and Prevention, National Center for Health Statistics. http://wonder.cdc.gov/ucd-icd10-expanded.html

Nisbett, R. E., & Wilson, T. D. (n.d.). The Halo Effect: Evidence for Unconscious Alteration of Judgments. *Journal of Personality and Social Psychology*, *35, 250–256*, 7. https://doi.org/10.1037/0022-3514.35.4.250

Peterson, C. R., & Beach, L. R. (1967). Man as an intuitive statistician. *Psychological Bulletin*, *68*(1), 29–46. https://doi.org/10.1037/h0024722

Pew Research Center. (2021, September 13). *Key facts about Americans and guns*. https://www.pewresearch.org/fact-tank/2021/09/13/key-facts-about-americans-and-guns/

Pierre, J. M. (2019). The psychology of guns: Risk, fear, and motivated reasoning. *Palgrave Communications*, *5*(1), Article 1. https://doi.org/10.1057/s41599-019-0373-z

Plassmann, F., & Whitley, J. (2003). Confirming "More Guns, Less Crime." *Stanford Law Review*, *55*(4), 1313–1369.

Porpora, D. V. & Sekalala, S. (2019). Truth, Communication, and Democracy. *International Journal of Communication*, *13, 18.*, 18.

Samimi, P., & Jenatabadi, H. S. (2014). Globalization and Economic Growth: Empirical Evidence on the Role of Complementarities. *PLOS ONE*, *9*(4), e87824. https://doi.org/10.1371/journal.pone.0087824

Smart, R., Morral, A. R., & Schell, T. L. (2021). The Magnitude and Sources of Disagreement Among Gun Policy Experts, Second Edition. *RAND Corporation*. https://doi.org/10.7249/RRA243-3

Soll, J. B., & Larrick, R. P. (2009). Strategies for revising judgment: How (and how well) people use others' opinions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*, 780–805. https://doi.org/10.1037/a0015145

Taber, C. S., Cann, D., & Kucsova, S. (2009). The Motivated Processing of Political Arguments. *Political Behavior*, *31*(2), 137–155.

Taber, C. S., & Lodge, M. (2012). Motivated Skepticism in the Evaluations of Political Beliefs. *Critical Review*, *24*(2), 157–184. https://doi.org/10.1080/08913811.2012.711019

Tappin, B. M., Pennycook, G., & Rand, D. G. (2020). Thinking clearly about causal inferences of politically motivated reasoning: Why paradigmatic study designs often undermine causal inference. *Current Opinion in Behavioral Sciences*, *34*, 81–87. https://doi.org/10.1016/j.cobeha.2020.01.003

Tenhundfeld, N. L., Parnes, J. E., Conner, B. T., & Witt, J. K. (2020). Development of a psychometrically valid gun attitude scale. *Current Psychology*, *39*(1), 279–286. https://doi.org/10.1007/s12144-017-9761-y

Therneau, T. (2023). *A Package for Survival Analysis in R*. https://CRAN.R-project.org/package=survival

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5(4), 297–323*, 27.

Watts, D. J., & Dodds, P. S. (2007). Influentials, networks, and public opinion formation. *Journal of Consumer Research*, *34*, 441–458. https://doi.org/10.1086/518527

Wintemute, G. J., Hemenway, D., Webster, D., Pierce, G., & Braga, A. A. (2010). Gun Shows and Gun Violence: Fatally Flawed Study Yields Misleading Results. *American Journal of Public Health*, *100*(10), 1856–1860. https://doi.org/10.2105/AJPH.2010.191916

Wojcik, S., Hilgard, S., Judd, N., Mocanu, D., Ragain, S., Hunzaker, M. B. F., Coleman, K., & Baxter, J. (2022). *Birdwatch: Crowd Wisdom and Bridging Algorithms can Inform Understanding and Reduce the Spread of Misinformation* (arXiv:2210.15723). arXiv. http://arxiv.org/abs/2210.15723

Wolfers, J., & Zitzewitz, E. (2004). Prediction Markets. *Journal of Economic Perspectives*, *18*(2), 107–126. https://doi.org/10.1257/0895330041371321

Wood, T., & Porter, E. (2019). The Elusive Backfire Effect: Mass Attitudes' Steadfast Factual Adherence. *Political Behavior*, *41*(1), 135–163. https://doi.org/10.1007/s11109-018-9443-y

Wyer, R. S., & Frey, D. (1983). The effects of feedback about self and others on the recall and judgments of feedback-relevant information. *Journal of Experimental Social Psychology*, *19*(6), 540–559. https://doi.org/10.1016/0022-1031(83)90015-X