Rational inference strategies and the genesis of polarization and extremism

Peter D. Kvam^{1,*}, Abhay Alaukik¹, Callie E. Mims², Arina Martemyanova¹, and Matthew Baldwin¹

¹University of Florida

²University of South Alabama

*Corresponding author: pkvam@ufl.edu

ABSTRACT

Polarization and extremism are often viewed as the product of psychological biases or social influences, yet they still occur in the absence of any bias or irrational thinking. We show that individual decision-makers implementing optimal dynamic decision strategies will become polarized, forming extreme views relative to the true information in their environment by virtue of how they sample new information. Extreme evidence enables decision makers to stop considering new information, whereas weak or moderate evidence is unlikely to trigger a decision and is thus under-sampled. We show that this information polarization effect arises empirically across choice domains including politically-charged, affect-rich and affect-poor, and simple perceptual decisions. However, this effect can be disincentivized by asking participants to make a judgment about the difference between two options (estimation) rather than deciding. We experimentally test this intervention by manipulating participants' inference goals (decision vs inference) in an information sampling task. We show that participants in the estimation condition collect more information, hold less extreme views, and are less polarized than those in the decision condition. Estimation goals therefore offer a theoretically-motivated intervention that could be used to alleviate polarization and extremism in situations where people traditionally intend to decide.

Scientific Reports: This is a preprint of an article published in *Scientific Reports*. The final authenticated version is available at doi.org/10.1038/s41598-022-11389-0

Author contributions: A.A., C.E.M., and A.M. created the stimuli and collected data, P.D.K. and A.A. analyzed the data., P.D.K. and M.B. conceived of the study and conditions in the main text, and all authors contributed to writing the paper.

1 Introduction

Attitudes toward crucial economic, environmental, and social policies are becoming increasingly polarized^{1,2}. Extreme, and polarized views that are exacerbated by social and motivational forces can lead to civil unrest and violence³. Meanwhile, polarization and extremism have begun to manifest more frequently beyond the political sphere, extending to divisions and conflict among juries, boardrooms, religious organizations, nonpartisan commissions⁴, sports, and even consumer product brands (Apple vs Google, car brands, or the “console wars”)⁵. Even in entirely mundane decisions about perceptual stimuli without any political charge, opinions can become polarized⁶. We often have little choice in the matter: Individuals' efforts and intentions to avoid forming extreme or polarized views can be rendered ineffective in the face of social and persuasive processes that exacerbate differences in opinion among groups⁷.

Polarization and extremism may not be so concerning if there were an easy solution to reversing their effects on cognition and behavior. Once it has set in, polarization distorts our view of the world and the way we consider new information, causing us to seek out confirming evidence^{8–10} both because confirming information is pleasurable and because it is felt to be more valuable than disconfirming information¹¹. Managing the dissonance that arises from exposure to counter-attitudinal information leads to greater polarization, as decision makers with opposing attitudes begin to deal in extremes. Common ground is lost as information is selected not to inform but rather to reinforce negative attitudes toward one's opponents^{12,13}. In concert with these confirmation biases, people also strive to disconfirm information that is not in line with their existing beliefs¹⁴. In terms of reasoning, this can be perfectly rational: Bayesian reasoning prescribes that we should not ignore our prior beliefs when forming new ones, even when processing counter-attitudinal information¹⁵. However, prior beliefs or expectations can affect how people interpret new information, resulting in them taking more extreme positions or forming more extreme beliefs in light of information that conflicts with their existing beliefs^{14,16–18}. People engage in motivated reasoning and belief revision driven

by a desire to form beliefs that align with others¹⁹ and the utilities associated with holding and maintaining those beliefs²⁰. This is exacerbated by the patterns of behavior on social networks, where communication and network dynamics lead to increases in polarization and the creation of echo chambers among users who agree with one another^{21–23}. Attempts to reverse the process of polarization can even backfire^{14, 16, 17}, as experimental findings have shown that groups of US Twitter users on opposite sides of the political spectrum became *more* polarized when exposed to opposing political views²⁴.

Considering how insidious polarization can be once it has taken root, it is critical to understand and prevent the process of polarization at its source. In this paper, we seek to explain how individual-level cognitive mechanisms can create polarization in the first place, examining how people gather information as they consider new situations. Specifically, we evaluate how *inference goals* interact with decisions about when to continue or stop sampling information when choosing among options^{25, 26}, and how different goals can create or ameliorate polarization at the genesis of new opinions or views. Computational modeling of these choice processes has revealed that individual-level decision dynamics are related to political beliefs and ideological preferences across a range of topics like authoritarianism and dogmatism, social dominance, and openness to opposing views more generally²⁷. Here, we go a step further to show that the very process of decision-making, even executed by a perfectly rational decision-maker, will create polarization and extreme views relative to the true information in the decision-maker's environment. That is, rational individual decision-makers who try to make accurate and efficient choices will still become polarized and adopt extremist views. And although we build on work suggesting that rational goals can create polarization in social communication and group environments^{28, 29}, observing polarization in our findings would suggest that rational individuals can form extreme and polarized views even in absence of communication or social pressure, in purely individual decisions, about relatively neutral and novel information ecologies.

If polarization is inevitable when making decisions, how can it be stopped? Fortunately, making a decision is not the only goal people can have when gathering new information. As opposed to determining which of two options is superior, one can evaluate or estimate their relative merit. Such a strategy incentivizes *precision* rather than *discrimination*, changing the conditions under which people stop sampling new information²⁶. As a result, changing the incentives or directions for a task to favor estimation may prevent people from forming polarized views in the first place. This shifts the target of polarization interventions to one of prevention, rather than treatment: Estimation incentives should keep people from ever becoming polarized or forming extreme views (relative to the true information “out there” in the world) in the first place. In the following sections, we outline the reasoning behind this kind of intervention and how it avoids polarization and extremism caused by decision goals, briefly note the results of a re-analysis of seven decision-making studies illustrating the prevalence of polarization, and introduce a new study where we manipulate the goals of information collection to reduce the extremeness and polarization of views people form.

1.1 The rational decision maker

Prior work often paints a picture of polarization as a phenomenon that arises from biases in information processing or belief updating³⁰, motivation^{8, 9}, ingroup-outgroup processes^{31, 32}, and media algorithms and design³³. Each of these explanations for polarization assumes that polarization emerges from bias or irrationality in the way decisions are made, while others explain how these beliefs can be propagated or exacerbated by rational reasoning based on already-polarized prior beliefs^{14, 16, 17, 27, 34–37}. While such biases or priors certainly perpetuate divisions between people and exacerbate the problem of extremism, they do not entirely explain the genesis of polarization and how similar people can wind up on opposite sides of an innocuous issue at the outset. In this paper, we suggest that polarization can be created on entirely novel issues from basic judgment and decision making processes themselves. Moreover, we argue that these processes can act independently of social forces, such as motivation or group-processes, and can lead to polarization even in the absence of initial biases or social interaction. To support this claim, we prove mathematically that decision-makers will be biased toward extreme views when implementing optimal choice strategies aimed at maximizing decision quality and efficiency. This proof is provided in its completeness in the supplementary material, but we summarize it here.

When a decision maker is presented with a choice between two (or more) options, they must make a compromise between investing time on that decision and making a choice so that they can move on. Both time and belief accuracy bestow utility – the utility of making good (accurate) decisions is clear, but making a quick decision allows a decision-maker to proceed to the next choice, or move on to a task that they enjoy more than the choice. If a decision maker focuses purely on collecting information, they will never be able to stop and make a choice, whereas if they focus entirely on decision speed then they will have at-chance accuracy and make generally poor choices. Under the general framework of random walk / diffusion processes that are commonly used to describe and explain choice behavior in humans and other animals^{38–40}, the optimal trade-off is achieved by setting a pre-defined *threshold* that corresponds to the level of accuracy that the decision maker desires, and then gathering information until that threshold is reached^{41, 42}. These models are based on the sequential probability ratio test^{42, 43}, where a decision maker tracks the balance of support between two options and uses this balance to determine when to make a choice. As they get information favoring option A, they shift up toward the “Choose A” response boundary, as shown in Figure

Figure 1. Illustration of how an unbiased distribution of information (left) becomes a polarized distribution (right) when fed through the decision process (middle). Because of an incentive to choose between options, moderate information – since it does not give credence to either options – gets discarded (middle)

1. Similarly, as they get information favoring option B, they shift down toward the “Choose B” response boundary. Once they cross one boundary or the other, the corresponding alternative is chosen (see the supplement for a full mathematical proof and description of simulations).

This strategy is specifically optimal for maximizing *reward rate*, or the amount of rewards that a decision-maker can accumulate per unit of time (seconds, minutes, days). There are slight modifications to the strategy, such as collapsing choice boundaries^{44,45}, that are made to adapt it to different environments, but all reward rate optimality models share the elements of a) tracking the balance of support, b) accumulating the relative balance to a threshold, and then c) deciding when the balance crosses a boundary. However, as we illustrate below, these three elements lead decision-makers to a distorted view of the evidence in their environment despite the decision process being optimal.

This unintended outcome occurs because a decision-maker’s goal is not to gather a representative sample of information, but to tip the balance of information far enough in one direction to conclude $A \succ B$ or $B \succ A$ (where $\succ$ denotes a preference order or belief). That is, choice tasks will bias information samples in the direction of the choice. On one hand, this is not so surprising, because we should expect people to have choice-consistent samples: Those who sample more of choice A information should choose option A. However, a deeper analysis reveals a surprising corollary to this expected pattern, namely, that the information sampled up to the point of choice is not representative of the true information in the world. Rather than mirroring the information in the environment (about choice A for instance), a decision-maker’s sampled information is more extreme^(25,26). In other words, the *goals* of the decision-makers influence the final sample of information that they collect, creating an unrepresentative sample of information despite no inherent biases.

One explanation for why this occurs is because moderate information is not conducive to achieving an agent’s decision goals (to stop and support $A \succ B$, or $B \succ A$), while extreme information makes it easy to stop and select one option over the other. These ideas are illustrated in Figure 1: Moderate information (pink) cannot tip the balance of evidence sufficiently toward one option or another to trigger a choice, while extreme information (red) is very likely to do so. As a result, people who have made a decision using a balance-of-evidence stopping rule will have a greater number of extreme information samples than the underlying distribution.

In addition to polarization, optimal decision strategies can create subgroups of highly polarized individuals whose information samples are at the extreme ends of the spectrum. When a person gathers a piece of strong information in their first sample, there is the chance that this is the only piece of information they need to consider to make a choice. As a result, some individuals are likely to construct their views and make their decisions based on a minimal sample of particularly extreme information. In the simulations in the supplementary material, we examine factors that exacerbate the problem of these “extremists” who gather very little information and yet hold much more extreme views than other decision makers. In general, these extremist decision makers tend to be implementing optimal strategies that place greater subjective value on time over new information relative to other decision makers, reflected in lower thresholds that generate fast but not necessarily accurate responses^{46,47}.

In the supplementary material, we provide a re-analysis of seven existing studies that allowed participants to make decisions by sampling information piece by piece over time. These studies show that the information participants consider is more extreme than the information naturally provided by their environment, lending credence to our prediction that decision-makers will become polarized. Furthermore, these studies show that the participants with the most extreme views are the same participants who sample the smallest amount of information, showing that extremism can arise even on simple perceptual and preferential choices. But the emergence of polarization and extremism is not inevitable. In the next section we describe one way that the decision process can be manipulated to incentivize the sampling of more representative information from the world.

1.2 Interventions to reduce polarization

The seven re-analyzed studies illustrate that polarization and extremism both arise frequently around even mundane choices (numbers of dots, which posters / foods are best), and it may seem inevitable in all decisions. Fortunately, decision-making is not the only goal that people can pursue when making inferences. Strategies that are non-optimal for reward rate maximization, but satisfy other criteria, can be useful for reducing polarization among groups of people. For example, a strategy elicited by incentivizing estimation accuracy – where a person is asked to assess *how much* the information in their environment favors one option relative to another – can lead participants to gather a distribution of information that is representative of the true information in the environment.

Some support for the efficacy of an estimation goal in reducing the sampling issues in decision making is provided by Coenen & Gureckis²⁶. They found that participants gathered largely representative samples while estimating relative strength of information, as opposed to a bimodal distribution of samples while choosing between two options. Model-based simulation results provide further evidence for its efficacy. As illustrated in Figure 2, an optimal strategy for making an estimate (minimizing response time for a desired level of error) requires a participant (i.e., the judge) to continue sampling new information until a desired level of precision is achieved. This process involves collecting new information and shifting their estimate over time (blue line) until the error of their estimate of the mean (shaded blue area can correspond to standard error or posterior variance of the mean) falls below a criterion level. In this case, extreme information is counterproductive to making precise estimates because it fails to reduce the variance of a sample, and thus the posterior variance of the mean. Conversely, information that is close to the current posterior mean results in maximal reduction of the posterior variance of the mean estimate. As a result, the judge is incentivized to gather more information to form an estimate after gathering extreme information, and incentivized to stop after gathering moderate information.

Figure 2. Illustration of a precision-based judgment process, where the goal is not to select one side over the other (as in Figure 1) but to meet a criterion level of estimation precision.

Following this optimal strategy, the distribution of sampled information across a population of decision makers (right side of Figure 2) will correspond closely to the input distribution (left side of Figure 2). Posing the task as an estimation problem can therefore be used as an effective manipulation that reduces polarization and extremism by bringing the evidence sample closer to the true information distribution. In this paper, we test the effectiveness of this intervention across multiple types of inference problems, encompassing both simple perceptual problems and more politically or affectively charged situations. In the next section, we introduce an experiment that examines these possibilities and explores the individual differences that drive or exacerbate polarization and extremism in some people.

2 Methods

All methods were carried out in accordance with relevant ethical guidelines and regulations. All experiments involving primary data were approved as exempt by the University of Florida Institutional Review Board (IRB202002176). These experiments were deemed to involve minimal risk to participants. Prior to completing the experiments, all participants were briefed on the study procedures and completed informed consent.

The experiments we present here were collected in two parts. An initial study collected data from 40 participants in each condition, separately for the choice and estimation conditions. We then collected data from a preregistered replication, randomly assigning an additional 110 participants to either the choice or estimation condition. However, the behavioral data from these two stages of data collection were nearly identical. We therefore present them altogether here for the sake of brevity. The data files from the first stage of data collection (Stage 1) and from the preregistered replication (All Data) are both provided on

the Open Science Framework at osf.io/jsb52, so readers can view and analyze the two data sets separately as they wish. The preregistration is available at osf.io/qfb6w.

Our initial re-analysis showed that choice tasks, by their very structure, lead to polarized and extreme information sampling, even when applying rational and optimal sampling strategies. Moreover, our own simulations as well as work by previous authors²⁶ suggest that estimation tasks could serve as a solution to this problem by reducing the extent to which people make decisions on the basis of extreme information. We further elucidate these processes in a new study that allows us to test whether polarization arises naturally in new environments with politically and affectively charged decisions. By comparing decision and estimation conditions, we evaluate whether estimation is a strategy that can alleviate polarization and extremism. The new study allows us to formalize our definitions of polarization and extremism (outlined below), and examine them as individual differences that can be related to relevant psychological traits, such as dogmatism, need for structure, and other demographic variables²⁷. In the sections to follow, we detail each of these measures and examine their relationship to information sampling and individual-level metrics of polarization and extremism.

The structure of choice and estimation tasks were matched closely: In both conditions, corresponding scenarios were worded almost identically, the layout of the window was kept identical, and the sampling action (clicking on a button to collect information about Policy A vs. B) was kept the same as well. The key difference between the two conditions was in the way participants responded to each scenario. In the choice condition, they were asked to decide between two courses of action ("Select Policy A or Policy B"). In the estimation condition, they were asked to estimate the relative expected effect of implementing the two courses of action on a relevant criterion (rate how good is Policy A vs Policy B). This allowed us to directly compare patterns of information sampled in the choice and estimation conditions for every question that we asked.

2.1 Task and items

In both conditions, participants were first instructed to imagine that they had encountered an alien civilization seeking their advice on several of their planet's pressing issues. The use of an "alien world" cover story is designed to minimize the influence of biases that participants might bring to each scenario and to make it clear to participants that they are gathering information with the goal of forming beliefs about an entirely new issue. The alien worlds paradigm has been commonly used to examine processes related to the evaluative information ecology⁴⁸, acquisition of morality⁴⁹, language learning⁵⁰, and reinforcement learning⁵¹.

Participants were informed that the aliens were considering two possibilities for each issue (i.e., "The aliens are building a new power plant that will generate energy from one of two materials – Material A and Material B") and participants were tasked with helping the aliens decide by sampling information about each of the options. Participants had access to a button that would generate the relevant information (e.g., "Each unit of Material A produces X MORE/LESS units of energy than each unit of Material B"). Each time participants clicked the button, the random variable X was drawn from a normal distribution with mean 0, $X \sim N(0, SD)$. This meant that the true information that participants could sample was normally distributed and (on average) did not favor one option over the other. Participants had the freedom to sample as much or as little as they wished before responding.

Crucially, participants in the choice condition were asked to choose the option which they thought was better (i.e., Material A or B), whereas participants in the estimation condition were asked to assess the degree to which one option was better than the other using a sliding scale (Figure 3) that encouraged participants to view the scenario as distinct from a binary choice. The actual choice or estimate that participants made was largely unimportant to the analyses we present below – instead, our aim was to evaluate the relative degree of polarization and extremism between the choice and estimation conditions by comparing patterns of information search in each scenario.

Participants in both conditions experienced a total of 80 scenarios: 40 in which the button enabled them to sample information from the world (i.e. gathering information using their "equipment" or "tools") and 40 in which the button allowed for sampling information from others (i.e. participants could consult with local alien experts to obtain relevant information). The scenarios were designed to incorporate a variety of issues that differed in terms of their subject matter as well as the degree to which they evoked affective involvement^{52,53}: the 80 questions included 20 affect-poor, 16 affect-rich, 22 politically charged, 8 risky, 10 investment, and 4 attention-check scenarios. A complete list of the choice scenarios, along with the data and analyses presented here, is provided on the Open Science Framework at osf.io/jsb52. Representative examples of scenarios under each category are shown in Table 1.

2.2 Participants

We recruited a total of 95 Prolific Academic workers for each condition for a total of 180 participants. This was composed of two stages of data collection: an initial sample of 40 participants in the choice and estimation conditions, assigned to conditions based on when they signed up to participate; and a preregistered replication study that randomly assigned 110 participants to either the choice or estimation condition. Our goal was to obtain data from approximately 85 participants in each condition, which is sufficient to identify relationships between individual-difference measures when they are present and to support the

The aliens are building a new power plant that will generate energy from one of two materials - Material A and Material B. Your task is to gather information about these materials using your equipment and recommend which material they should choose.

Measure Energy Output

Measurements show that material A generates about 38 kWh MORE energy than material B.
Measurements show that material A generates about 11 kWh MORE energy than material B.
Measurements show that material A generates about 12 kWh LESS energy than material B.

Once you are ready to choose, select one of these options:

Recommend Material A

Recommend Material B

The aliens are building a new power plant that will generate energy from one of two materials - Material A and Material B. Your task is to gather information about these materials using your equipment and estimate the difference in output between Material A and Material B.

Measure Energy Output

Measurements show that material A generates about 6 kWh MORE energy than material B.
Measurements show that material A generates about 2 kWh MORE energy than material B.
Measurements show that material A generates about 30 kWh LESS energy than material B.

Once you are ready, enter your estimate using the scale below:

Material A will generate 50kWh LESS than Material B Material A will generate 50 kWh MORE than Material B

-50 -40 -30 -20 -10 0 10 20 30 40 50

Estimated difference in energy output:

Figure 3. The scenarios in both tasks are worded very closely. The key difference lies in the task: after gathering information by clicking the button, participants either choose one of two options (left), or estimate how much better one option is than the other (right; the scale ranges from -2.5 to +2.5 SD of the true information distribution).

Table 1. Representative examples of issues under each category of scenario.

Category	Representative example
Affect-Rich	Pursue gene therapy or immunotherapy to save aliens from a deadly virus.
Affect-Poor	Mining for the mineral Xenon in Location A or B.
Politically Charged	Allowing immigrant aliens into the city or not.
Risky	Settling in city A or B based on earthquake frequency.
Investment	Investing in flood prevention infrastructure now or spend money later on flood response.

null hypothesis with a Bayes factor of 3 or more (moderate evidence) when there is no relationship between measures. A formal power analysis is provided in the supplement.

Ten participants from the choice condition and eleven from the estimation condition were removed because they failed attention checks or failed to respond at all on multiple trials. Hence, we retained responses from 85 individuals in the choice condition and 84 in the estimation condition, very close to our goal of 85 participants per condition. Participants in the choice condition were 54 women (30 men, 1 non-binary / other) with a mean age of 33.62 ($SD = 11.63$). For the estimation condition, participants were 45 women (37 men, 2 non-binary / other) with an average age of 36.59 ($SD = 12.71$). All participants indicated that they were from the United States and fluent English speakers. Participants were paid \$10 per hour for their participation in the experiment.

Our sample size was intended to allow us 80% “power” – in this case, the ability to conclude in favor of the null using a Savage-Dickey Bayes factor analysis⁵⁴ when there was no relationship between variables in a regression. Moreover, we collected responses to 80 trials for each participant, which results in approximately 13,500 data points in total to estimate the effects. Such a large set of data helps make our parameter estimates of both group-level and individual-level characteristics robust and reliable. Moreover, most of our analyses involve hierarchical (mixed effect) Bayesian models, which are typically stricter than classical statistics under uninformative priors⁵⁵. In the results, we also provide 95 percent credible / highest density intervals [HDIs] for our point estimates, which indicate the 95% posterior most likely estimates of a particular parameter or statistic.

2.3 Individual difference analyses

Similar to the re-analysis of the seven empirical studies in the previous section, we compared the observed distribution of information collected by participants against the expected distribution of an unbiased, representative sample that we would expect from a random sample of information. These distributions were compared in two main ways. First, we computed

Figure 4. Diagram of the measures of polarization (low = left, high = right) and extremism (low = top, high = bottom), how different distributions of information might exemplify each of these properties, and how KLD and variance ratio (VarRatio) describe / quantify these properties. The dotted black line corresponds to the expected distribution, a normal distribution centered at zero, while the blue line corresponds to the observed distribution with high or low polarization / extremeness.

the difference between the distributions using an information-theoretic metric, the Kullback-Leibler divergence [KLD]. This quantifies the difference between the expected probability density of samples (a unimodal, normal distribution) and the observed density of samples. It is used as our measure of polarization, as it quantifies the degree to which the samples gathered by participants diverge from an expected unimodal distribution.

Second, we also quantify the extremeness of the sampled information by evaluating its variance, divided by the true variance of the pool of information participants are sampling from (i.e., the variance of a representative sample). This allows us to assess the degree to which participants' sampling strategies led them to draw samples that are more (or less) extreme than the true information provided by their environment. Note that this measure of extremeness is often correlated with, but distinct from, our measure of polarization. As shown in Figure 4, a distribution of information can be highly polarized / bimodal without high variance, and likewise it can have high variance without being bimodal or polarized.

Putting it all together, we formally define polarization and extremism in these studies as follows:

- **Polarization:** Polarization is the extent of divergence between the distribution of the information collected and the true (expected) information distribution. We operationalize it as the Kullback-Liebler Divergence [KLD], which provides a measure of the overall difference between two probability distributions that we can use to compare observed and expected information distributions:

$$KLD[p(y)||p(x)] = \int_i^N p(\mathbf{y}_i) \log_n \frac{p(\mathbf{y}_i)}{p(\mathbf{x}_i)} \quad (1)$$

where $p(\mathbf{y}_i)$ describes the probability density of a standard normal distribution and $p(\mathbf{x}_i)$ is a nonparametrically estimated probability density, generated using a Gaussian kernel density estimator with optimal bandwidth⁵⁶, of sampled data. The indices $i : N$ represent each of points on the probability distributions where their divergence is calculated.

- **Extremism:** The ratio of the variance of the observed sample to the variance of a representative sample of information

Figure 5. Distribution of information collected by participants (blue histogram) in the choice (left) and estimation (right) conditions, compared against the expected distribution that would result from unbiased, representative sampling (orange line).

from that trial – termed *variance ratio* for the purposes of this study. If the overall variance ratio in the condition (i.e., condition-level) is above 1, the information sampled in the condition will be deemed extreme.

Both of these metrics can be applied at the level of a trial, participant, or entire condition: We simply evaluate the KLD or variance of a single trial relative to the expected distribution on that trial, all trials from a participant relative to the expected distribution for that participant, or all trials from all participants in a condition, respectively. In all cases, we compare the observed to expected distributions *conditional on how many samples were drawn*. This allows us to disentangle the number of samples drawn from the shape of the distribution of samples they drew. It also allows us to examine the number of samples as a separate outcome variable. The number of samples (the number of clicks on the “draw another sample” button in Figure 3) served as a simple and direct measure of how much information people chose to gather before making their choice or estimate. We also measured the length of time participants spent on each trial in seconds, although this data was not analyzed.

For the extremism analysis / variance ratio, we can estimate uncertainty in the variance parameter simply by fitting it using Bayesian MCMC sampling methods. However, the KLD is a point estimate and thus does not convey uncertainty, making it hard to directly compare across conditions. To estimate a 95% HDI on the KLD for decision and estimation tasks, we instead simulated 10,000 artificial data sets with the same characteristics (number of trials, number of samples per trial) as the observed data. However, these artificial data sets were sampled from the true underlying distribution of information from the stimulus, i.e., they did not use the same stopping rules as participants. In other words, they did not have any sampling biases that our participants might have had. This yielded 10,000 “samples” from the true distribution against which we could compare the observed results. We computed the mean difference between observed and expected KLDs by subtracting the mean of the KLDs across all 10,000 artificial data sets from the KLD of the observed data set. Then, we constructed the 95% HDI on the difference between the observed and expected distribution by computing the 95% highest density interval of differences between observed KLD and each of the 10,000 artificial data set KLDs.

At the conclusion of the study, participants completed self-report measures of dogmatism⁵⁷ and personal need for structure⁵⁸ and reported their political ideology and basic demographic information. We included dogmatism in our analysis because previous work has connected it to reduced information sampling under uncertainty⁵⁹. Need for structure was included because need for closure, of which the need for structure appears to be a more psychometrically desirable variant⁵⁸, has previously been connected to the threshold parameter in dynamic decision models⁶⁰. We report the relationships between all of these individual differences – extremism / variance ratio, polarization / KLD, information sampling / number of clicks, dogmatism, and need for structure – in the results.

3 Results

All the analyses were conducted in R⁶¹, JASP⁶², and MATLAB. Data were analyzed with Bayesian estimation of the posterior using uninformed priors and hierarchical / mixed models, where appropriate. We implemented MCMC sampling in JAGS⁶³ to generate all statistics of interest (such as variance ratios and correlations). We report mean estimates and their corresponding 95 percent HDI⁵⁵. Code for these analyses are provided on the Open Science Framework at (osf.io/jsb52), and we expand on their descriptions in the supplementary materials.

3.1 Polarization and Extremism

We first compared the distributions of information collected by participants against the distributions we would expect from a random, representative sample across trials. These distributions can be seen in Figure 5. The blue histograms represent the distributions of information sampled by participants, while the orange line represents the true (standard normal) distribution from which the samples were drawn. The difference is immediately striking – as in the re-analyzed studies and the model prediction shown in Figure 1, there is a “dip” in the middle of the distribution of information collected by participants in the choice condition. This bimodal distribution of information shows that people are under-representing moderate information when sampling.

By comparison, the estimation condition shows little to no substantial deviation from the expected distribution, indicating that the estimation intervention appears to evoke more representative sampling than the choice condition. That is, moderate information was considered equally in making estimations, resulting in a typical bell-shaped distribution of information that reflected the true underlying information in the environment.

Our quantitative measure of polarization (KLD) also aligned with these visual observations: the expected KLD in the decision condition was approximately 0.12, but the observed KLD was 1.00 ($M_{diff} = 0.88$, 95% HDI = [0.80, 0.95]). The expected KLD in the estimation condition was approximately 0.12, whereas the observed KLD was 0.52 ($M_{diff} = 0.40$, 95% HDI = [0.32, 0.47]). Thus, the deviation from the true distribution – and thus the amount of polarization, as defined by KLD – was more than twice as large in the decision condition compared to the estimation condition. KLD scores in the estimation condition were on average .26 standard deviations lower (95% HDI = [−.30, −.23]) than KLD scores in the choice condition, corresponding to a medium-sized effect⁶⁴. A model comparison revealed that the regression yielded extremely strong support for inclusion of condition as a factor predicting KLD ($\log(BF) = 358.30$).¹

Turning to extremism metrics, we found the condition-level variance ratio for the choice condition to be $M = 1.18$ (95% HDI = [1.11, 1.27]), where a value of 1.00 indicates that the observed sample is equally as extreme as a representative sample. Conversely, in the estimation condition, the variance ratio was centered right around the expected variance, $M = 1.01$ (95% HDI = [0.96, 1.09]). These HDIs are entirely non-overlapping, meaning that we can be very confident that the variance in the choice condition is higher than the variance in the estimation condition, signaling that the estimation task reduced the extremeness of information participants gathered.

We also evaluated patterns of extremeness by calculating correlations between the variance of samples people collected with the total number of times they chose to sample information (i.e., number of clicks). The extremeness of information participants collected was credibly correlated with the number of times people sampled information in each trial, both in the choice condition ($r = -0.44$; HDI = [−0.47, −0.41]; $\ln(BF) = 296.16$), as well as in the estimation condition ($r = -0.48$; HDI = [−0.51, −0.46]; $\ln(BF) = 397.94$). That is to say, the more people sampled extreme information, the fewer times they clicked to receive new information *within both conditions*. The amount (or lack of) of information collected is therefore predictive of extremism above and beyond choice-based polarization, suggesting that manipulations increasing information sampling are likely to be effective in both choice and estimation tasks.

3.2 Information search across conditions

We hypothesized that one of the reasons for extremism and polarization is that some participants tend to gather less information than others (lower thresholds), and less information typically leads to more volatile / extreme samples of evidence. It is possible that the choice / estimation manipulation, or the question types, led participants to seek out more or less information. This can be examined by looking at the number of pieces of information that participants collected in each condition, quantified as the number of times they clicked on the button to sample more information before making their choice or estimate. A comparison between the different question types, as well as between choice and estimation conditions, on the number of pieces of information sampled is shown in Figure 6.

To formally quantify whether sampling differed between choice and estimation conditions, we examined the number of samples that participants gathered in each condition. We ran a Bayesian Linear Regression in JASP⁶² with default / standard priors^{67,68}, which provided strong evidence that condition should be included as a predictor of the number of clicks ($\ln(BF) = 13.12$). In the estimation condition, number of clicks increased on average by 0.32 clicks relative to the choice condition (95% HDI = [0.21, 0.44]). Therefore, part of the way that the estimation condition reduces extremism is by encouraging participants to seek out more information, resulting in more informed judgments.

Similarly, a greater number of clicks led to less polarization, reducing the KLD by an average of 0.06 standard deviations per click in the choice condition (95% HDI = [−0.07, −0.05]; $\ln(BF) = 125.29$) and by an average of 0.09 standard deviations per click in the estimation condition (95% HDI = [−0.09, −0.08], $\ln(BF) = 284.67$). Gathering more information is a clear-cut

¹Note that we report natural log Bayes factors here and elsewhere in the text, as often the Bayes factors on raw odds scales are approach infinite values. A Bayes factor of 30 on a raw scale (+3.40 on a natural log scale) indicates extremely strong support for an effect / inclusion of a model parameter, while a Bayes factor of 1/30 on a raw scale (−3.40 on a natural log scale) indicates strong evidence against an effect / inclusion of a model parameter^{65,66}.

Figure 6. Mean number of samples drawn (error bars indicate standard error) for each question type in the choice (blue) and estimation (yellow) condition.

strategy for reducing both polarization and extremism, which is why the increase in information search in the estimation condition was so effective at reducing both effects.

To examine whether the type of inference domain affected sampling behavior, we tested whether condition-level manipulations impacted information sampling by looking at the effect of question type (affect poor, affect rich, etc) on number of clicks. We ran a Bayesian ANOVA in JASP^{62,69} which indicated that question type had no effect on number of clicks ($BF(H_0 > H_1) = 270.66$, indicating strong support for the null hypothesis). There were also no substantial differences on question type in predicting individual-level KLDs ($BF(H_0 > H_1) = 1237.94$, indicating strong support for the null hypothesis) so we do not examine question type further.

Table 2. Estimates of the effects of Dogmatism, Personal Need for Structure Reaction to Lack of Structure subscale (PNS-RLS), and Personal Need for Structure Desire for Structure subscale (PNS-DFS) on polarization (K-L divergence, bottom) and information search (number of clicks on the sampling button on each trial, top).

Information sampling			
Term	Estimate	SE	95% HDI
Intercept	4.14	0.23	[3.70, 4.60]
Dogmatism	-0.01	0.00	[-0.01, -0.00]
PNS-DFS	-0.14	0.01	[-0.16, -0.12]
PNS-RLS	0.23	0.01	[0.20, 0.26]
Polarization			
Term	Estimate	SE	95% HDI
Intercept	47.22	2.90	[41.43, 52.99]
Dogmatism	-0.09	0.05	[-0.18, -0.00]
PNS-DFS	0.55	0.26	[0.04, 1.08]
PNS-RLS	-0.70	0.36	[-1.42, -0.01]

3.3 Individual Differences

In addition to the differences between conditions, we hypothesized that information search and polarization/extremism should be predicted by personality factors like dogmatism (DOG) and personal need for structure (PNS), which is comprised of the desire for structure (DFS) and reaction to lack of structure (RLS). As summarized in Table 2, the model predicting information search indicated a negative effect of dogmatism, positive effect of the personal need for structure - desire for structure subscale, and negative effect of the reaction to lack of structure subscale. This aligns with recent findings showing that dogmatism is related to decreased information search under uncertainty^{27,59}. Interestingly, the sub-scales of Personal Need for Structure predicted information search in opposite directions. It appears that a desire for structure leads participants to search for less

information, perhaps because making a quick decision satisfies the desire to resolve uncertainty and reach a conclusion, while the discomfort associated with having ill-informed views based on little evidence (reaction to lack of structure) drives people to sample more.

These individual differences also predicted individual-level polarization. This was tested by again using a Bayesian linear regression to assess the credibility of individual differences as predictors of participant-level (standardized) KLD. As with information search, the regression suggested that all three measures predicted KLD: dogmatism led to decreased polarization, desire for structure led to greater polarization, and reaction to lack of structure led to less polarization.

The KLD results are particularly noteworthy because dogmatism predicts behavior in the opposite direction than we expected: Greater dogmatism predicted less polarization in resulting sampling behavior (lower KLD), while the PNS scales align with the information search results. So while dogmatism results in reduced information search, this reduction in information sampling does not translate directly to less representative views. One possibility is that highly dogmatic individuals hit a “sampling floor” and sampled only one piece of information before deciding. This would result in a low KLD, because lone samples will not favor extreme information more than moderate information and will, in fact, constitute representative simple random samples. By setting very low thresholds, dogmatic individuals therefore shift back toward a representative sample. As a result, they are ill-informed but no more polarized than a random sample of information from the environment would expect.

4 Discussion

Our studies illustrate that regardless of the domain – spanning perceptual choices about dots or colors, preferential choices about food or posters (as revealed in the reanalysis of seven existing studies), neutral social policies, or politically and affectively charged decisions (as revealed in the empirical studies reported here) – the process of decision making generates polarization and extremism during the information sampling process. These were scenarios where there was no social interaction among participants, and where there should be little motivation to engage in biased information processing. While a tremendous volume of work has gone into understanding motivated reasoning⁷⁰, communication dynamics and social networks^{21,28,29}, and the social forces that shape and exacerbate polarization^{31,32,71,72}, it has focused less on individual-level information search mechanisms that might create polarization before beliefs impacted by social influence. Our work suggests that the root causes of polarization may be even deeper than previously thought, arising from completely rational, individual-level decision strategies.

As we show, polarized and extreme individuals may be those who are merely reacting rationally to choice incentives, and are not necessarily intentionally malicious actors. Based on the model simulations, we should expect that extremists will show up most frequently when there is time pressure to make choices and when decision makers begin the choice process with a biased position. Nevertheless, rational decision making processes – resulting from dual pressures on decision quality and time – will create polarized groups of decision makers and generate a subset of uninformed individuals who hold extreme views.

A major concern regarding these extremists is the degree of influence that they can carry in cultural discourse and the formation of public opinion. Decision makers with small, extreme samples of information are still confident in their views^{73,74}, possibly due to a belief in the “law of small numbers” where people think even very small samples will possess the same statistical properties as large samples⁷⁵. This suggests that extremists will have strong convictions and be willing to spread their views, despite the information they have gathered possessing relatively poor reliability. This is exacerbated by several other established phenomena. First, those with the most extreme positions carry a disproportionately greater sway over group discourse because their outlying beliefs are more robust to change and can thus “pull” other people toward them⁷⁶. Second, extremists will be the first to share the information they have – by virtue of taking less time to make decisions, they are free to share their views before those who have taken more time and carefully constructed their beliefs with large samples of information. Extremists will therefore be the first to influence their peers via word of mouth and social networks⁷⁷, provide reviews of new products that influence many subsequent buyer decisions⁷⁸, and spread hateful views that carry quickly through social networks⁷⁹. Social media algorithms base recommendations on early posters, meaning that their “hot takes” will carry greater weight in determining subsequent users’ views and activity as well^{77,80}.

As a result, we arrive at a scenario where the least-informed (fewest samples of information), most dogmatic, and most impatient (lowest-threshold) individuals in a population will have the greatest influence on public opinion. The high profile of extremists can explain why perceived political polarization is exaggerated—people tend to think that the political left and right in the US are far more polarized than they are⁸¹.

4.1 Alleviating polarization & extremism

The estimation incentives we examined in the new study suggest that polarization and extremism can be reduced at the outset by changing the optimal sampling strategy. Simply asking people to give their best estimate of the difference between options, as opposed to deciding between them, drives a more natural information sampling approach that reduces polarization and extremism while promoting greater information search. Successful implementation of this type of intervention provides avenues

for reducing the divides between groups of people, and potentially can be brought to bear in designing simple interventions or judgment architectures to reduce these social problems⁸².

In the political sphere, our results suggest that binary or multiple-choice voting – requiring voters to decide in favor of only Candidate A or Candidate B (or C, D, E, etc) – may be partly responsible for the current state of political polarization. Fortunately, the estimation results suggest a remedy for this particular component of polarization. Specifically, it may be possible to alleviate political polarization through approaches like cardinal voting⁸³, which incentivizes precise ratings (estimation) as opposed to forcing voters to select one option over another. Based on our estimation results, there is reason to hope that this could encourage greater information search among voters as well as more representative sampling. Certainly future work ought to evaluate the practicality and efficacy of different voting interventions like these that could reduce polarization on social issues.

An intriguing possibility that is raised by this work is that many types of information search may yield polarized samples. This may apply beyond literal information search – as we presented here – into domains like information search from memory, which is thought to leverage analogous mechanisms to external information search^{38,84}. If search through memory is terminated using the same rules as diffusion or random walk models, then a person's judgment goals at the time of measurement (when they are asked to make a choice or estimate) may determine what information they consider. Similarly, it may be possible to override samples of biased information that have been collected by changing the communication process. Asking a polarized individual to share information to someone who's goal it is to *estimate* may incentivize sharing of representative information from memory. It may therefore be possible to alleviate polarization and extremism by manipulating inference and communication goals in more applied situations like the ballot box, boardroom, or social media feed. To the extent that the interventions we design here are useful for manipulating information search in these domains, they are promising for reducing polarization and extremism.

Conclusions

In a decision scenario that incentivizes a trade-off between time and decision quality, rational decision makers will sample more polarized and extreme information than their environment naturally provides. However, we have also identified a potential solution that is as simple as changing decision goals: Polarization and extremism disappear when people are asked to estimate some relative quantity or preferability instead of choosing between options. This approach to making inferences encourages information search while simultaneously reducing the biases inherent to decision-making. Mathematical models, simulations, re-analysis of data from seven empirical studies, as well as results from a new experiment point to the advantages of estimation tasks over choice/decision tasks if one wants to reduce polarization and extremism. Future work will determine how widely and effectively this type of intervention can be implemented to reduce polarization and extremism outside the laboratory, but we hope that this work sheds light on one route to responsibly construct inference environments to create more common ground that brings people together instead of driving them apart.

Availability of Data and Materials

The datasets generated and analysed during the current study are available in the Open Science Framework repository, osf.io/jsb52.

References

1. Baldassarri, D. & Gelman, A. Partisans without constraint: Political polarization and trends in american public opinion. *Am. J. Sociol.* **114**, 408–446 (2008).
2. Wilson, A. E., Parker, V. & Feinberg, M. Polarization in the contemporary political and media landscape. *Curr. Opin. Behav. Sci.* **34**, 223–228 (2020).
3. Berger, J. M. *Extremism* (MIT Press, 2018).
4. Sunstein, C. R. The law of group polarization. *Univ. Chic. Law Sch. John M. Olin Law & Econ. Work. Pap.* (1999).
5. Shi, Y., Mast, K., Weber, I., Kellum, A. & Macy, M. Cultural fault lines and political polarization. In *Proceedings of the 2017 ACM on Web Science Conference*, 213–217 (2017).
6. Baron, R. S. & Roper, G. Reaffirmation of social comparison views of choice shifts: Averaging and extremity effects in an autokinetic situation. *J. Pers. Soc. Psychol.* **33**, 521 (1976).
7. Isenberg, D. J. Group polarization: A critical review and meta-analysis. *J. personality social psychology* **50**, 1141 (1986).
8. Festinger, L. *A theory of cognitive dissonance*, vol. 2 (Stanford university press, 1957).
9. Nickerson, R. S. Confirmation bias: A ubiquitous phenomenon in many guises. *Rev. Gen. Psychol.* **2**, 175–220 (1998).

10. Westerwick, A., Johnson, B. K. & Knobloch-Westerwick, S. Confirmation biases in selective exposure to political online information: Source bias vs. content bias. *Commun. Monogr.* **84**, 343–364 (2017).
11. Lord, C. G., Ross, L. & Lepper, M. R. Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *J. personality social psychology* **37**, 2098 (1979).
12. Iyengar, S., Sood, G. & Lelkes, Y. Affect, not ideology: A social identity perspective on polarization. *Public Opin. Q.* **76**, 405–431 (2012).
13. Kuhn, D. & Lao, J. Effects of evidence on attitudes: Is polarization the norm? *Psychol. Sci.* **7**, 115–120 (1996).
14. Gershman, S. J. How to never be wrong. *Psychon. bulletin & review* **26**, 13–28 (2019).
15. Taber, C. S., Cann, D. & Kucsova, S. The motivated processing of political arguments. *Polit. Behav.* **31**, 137–155 (2009).
16. Cook, J. & Lewandowsky, S. Rational irrationality: Modeling climate change belief polarization using bayesian networks. *Top. cognitive science* **8**, 160–179 (2016).
17. Dixit, A. K. & Weibull, J. W. Political polarization. *Proc. Natl. Acad. sciences* **104**, 7351–7356 (2007).
18. Botvinik-Nezer, R., Jones, M. & Wager, T. Fraud beliefs following the 2020 us presidential election: A belief systems analysis. *PsyArXiv* (2021).
19. Wheeler, N. E. *et al.* Ideology and predictive processing: coordination, bias, and polarization in socially constrained error minimization. *Curr. Opin. Behav. Sci.* **34**, 192–198 (2020).
20. Zmigrod, L. Mental computations of ideological choice and conviction: The utility of integrating psycho-economics and bayesian models of belief. *Psychol. Inq.* (in press).
21. Kashima, Y., Perfors, A., Ferdinand, V. & Pattenden, E. Ideology, communication and polarization. *Philos. Transactions Royal Soc. B* **376**, 20200133 (2021).
22. Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A. & Van Bavel, J. J. Emotion shapes the diffusion of moralized content in social networks. *Proc. Natl. Acad. Sci.* **114**, 7313–7318 (2017).
23. Crockett, M. J. Moral outrage in the digital age. *Nat. Hum. Behav.* **1**, 769–771 (2017).
24. Bail, C. A. *et al.* Exposure to opposing views on social media can increase political polarization. *Proc. Natl. Acad. Sci.* **115**, 9216–9221 (2018).
25. Coenen, A. & Gureckis, T. The distorting effect of deciding to stop sampling. In Papafragou, A., Grodner, D., Mirman, D. & Trueswell, J. (eds.) *Proceedings of the 37th Annual Meeting of the Cognitive Science Society*, 2819–2824 (Cognitive Science Society, 2016).
26. Coenen, A. & Gureckis, T. The distorting effects of deciding to stop sampling information. *PsyArXiv* (2021).
27. Zmigrod, L., Eisenberg, I. W., Bissett, P. G., Robbins, T. W. & Poldrack, R. A. The cognitive and perceptual correlates of ideological attitudes: a data-driven approach. *Philos. Transactions Royal Soc. B* **376**, 20200424 (2021).
28. Burnstein, E. & Vinokur, A. Testing two classes of theories about group induced shifts in individual choice. *J. Exp. Soc. Psychol.* **9**, 123–137 (1973).
29. Burnstein, E. & Vinokur, A. Persuasive argumentation and social comparison as determinants of attitude polarization. *J. Exp. Soc. Psychol.* **13**, 315–332 (1977).
30. DeKay, M. L. Predecisional information distortion and the self-fulfilling prophecy of early preferences in choice. *Curr. Dir. Psychol. Sci.* **24**, 405–411 (2015).
31. Hogg, M. A. & Turner, J. C. Interpersonal attraction, social identification and psychological group formation. *Eur. J. Soc. Psychol.* **15**, 51–66 (1985).
32. Suhay, E. Explaining Group Influence: The Role of Identity and Emotion in Political Conformity and Polarization. *Polit. Behav.* **37**, 221–251, DOI: 10.1007/s11109-014-9269-1 (2015).
33. Benkler, Y., Faris, R. & Roberts, H. *Network propaganda: Manipulation, disinformation, and radicalization in American politics* (Oxford University Press, 2018).
34. Olsson, E. J. A bayesian simulation model of group deliberation and polarization. In *Bayesian argumentation*, 113–133 (Springer, 2013).
35. Jern, A., Chang, K.-M. K. & Kemp, C. Belief polarization is not always irrational. *Psychol. review* **121**, 206 (2014).
36. Pallavicini, J., Hallsson, B. & Kappel, K. Polarization in groups of bayesian agents. *Synthese* **198**, 1–55 (2021).

37. Singer, D. J. *et al.* Rational social and political polarization. *Philos. Stud.* **176**, 2243–2267 (2019).
38. Ratcliff, R., Smith, P. L., Brown, S. D. & McKoon, G. Diffusion decision model: Current issues and history. *Trends Cogn. Sci.* **20**, 260–281 (2016).
39. Busemeyer, J. R., Gluth, S., Rieskamp, J. & Turner, B. M. Cognitive and neural bases of multi-attribute, multi-alternative, value-based decisions. *Trends Cogn. Sci.* **23**, 251–263 (2019).
40. Brunton, B. W., Botvinick, M. M. & Brody, C. D. Rats and humans can optimally accumulate evidence for decision-making. *Science* **340**, 95–98 (2013).
41. Bogacz, R., Brown, E., Moehlis, J., Holmes, P. & Cohen, J. D. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* **113**, 700–765, DOI: 10.1037/0033-295X.113.4.700 (2006).
42. Edwards, W. Optimal strategies for seeking information: Models for statistics, choice reaction times, and human information processing. *J. Math. Psychol.* **2**, 312–329, DOI: 10.1016/0022-2496(65)90007-6 (1965).
43. Wald, A. & Wolfowitz, J. Bayes solutions of sequential decision problems. *Proc. Natl. Acad. Sci. United States Am.* **35**, 99–102 (1949).
44. Tajima, S., Drugowitsch, J. & Pouget, A. Optimal policy for value-based decision-making. *Nat. communications* **7**, 1–12 (2016).
45. Boehm, U., van Maanen, L., Evans, N. J., Brown, S. D. & Wagenmakers, E.-J. A theoretical analysis of the reward rate optimality of collapsing decision criteria. *Attention, Perception, & Psychophys.* **82**, 1520–1534 (2020).
46. Heitz, R. P. & Schall, J. D. Neural mechanisms of speed-accuracy tradeoff. *Neuron* **76**, 616–628, DOI: 10.1016/j.neuron.2012.08.030 (2012).
47. Wickelgren, W. A. Speed-accuracy tradeoff and information processing dynamics. *Acta Psychol.* **41**, 67–85, DOI: 10.1016/0001-6918(77)90012-9 (1977).
48. Alves, H., Koch, A. & Unkelbach, C. A cognitive-ecological explanation of intergroup biases. *Psychol. Sci.* **29**, 1126–1133 (2018).
49. Rottman, J. & Kelemen, D. Aliens behaving badly: Children’s acquisition of novel purity-based morals. *Cognition* **124**, 356–360 (2012).
50. Gupta, P. *et al.* Space aliens and nonwords: Stimuli for investigating the learning of novel word-meaning pairs. *Behav. Res. Methods, Instruments, & Comput.* **36**, 599–603 (2004).
51. Decker, J. H., Otto, A. R., Daw, N. D. & Hartley, C. A. From creatures of habit to goal-directed learners: Tracking the developmental emergence of model-based reinforcement learning. *Psychol. science* **27**, 848–858 (2016).
52. Zajonc, R. B. & Markus, H. Affective and cognitive factors in preferences. *J. consumer research* **9**, 123–131 (1982).
53. Pachur, T., Hertwig, R. & Wolkewitz, R. The affect gap in risky choice: Affect-rich outcomes attenuate attention to probability information. *Decision* **1**, 64 (2014).
54. Wagenmakers, E. J., Lodewyckx, T., Kuriyal, H. & Grasman, R. Bayesian hypothesis testing for psychologists: A tutorial on the Savage-Dickey method. *Cogn. Psychol.* **60**, 158–189, DOI: 10.1016/j.cogpsych.2009.12.001 (2010).
55. Kruschke, J. K. *Doing Bayesian Data Analysis: A Tutorial with R, JAGS, and STAN* (Academic Press, 2014).
56. Bowman, A. W. & Azzalini, A. *Applied smoothing techniques for data analysis: the kernel approach with S-Plus illustrations*, vol. 18 (OUP Oxford, 1997).
57. Altemeyer, B. Dogmatic behavior among students: Testing a new measure of dogmatism. *The J. Soc. Psychol.* **142**, 713–721 (2002).
58. Neuberg, S. L. & Newsom, J. T. Personal need for structure: Individual differences in the desire for simpler structure. *J. Pers. Soc. Psychol.* **65**, 113 (1993).
59. Schulz, L., Rollwage, M., Dolan, R. J. & Fleming, S. M. Dogmatism manifests in lowered information search under uncertainty. *Proc. Natl. Acad. Sci.* **117**, 31527–31534 (2020).
60. Evans, N. J., Rae, B., Bushmakin, M., Rubin, M. & Brown, S. D. Need for closure is associated with urgency in perceptual decision-making. *Mem. & cognition* **45**, 1193–1205 (2017).
61. Team, R. C. R: A language and environment for statistical computing (2013).

- 541 **62.** JASP Team. JASP (Version 0.14.1)[Computer software] (2020).
- 542 **63.** Plummer, M. JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the*
543 *3rd International Workshop on Distributed Statistical Computing*, vol. 124, 10 (Vienna, Austria., 2003).
- 544 **64.** Cohen, J. *Statistical power analysis for the behavioral sciences* (Academic press, 2013).
- 545 **65.** Jeffreys, H. *The Theory of Probability* (OUP Oxford, 1961).
- 546 **66.** Raftery, a. E. Bayesian model selection in social research (with discussion by Andrew Gelman & Donald B. Rubin, and
547 Robert M. Hauser, and a rejoinder). *Sociol. Methodol.* 111–196 (1995).
- 548 **67.** Clyde, M. A., Ghosh, J. & Littman, M. L. Bayesian adaptive sampling for variable selection and model averaging. *J.*
549 *Comput. Graph. Stat.* **20**, 80–101 (2011).
- 550 **68.** Consonni, G., Fouskakis, D., Liseo, B. & Ntzoufras, I. Prior distributions for objective bayesian analysis. *Bayesian*
551 *Analysis* **13**, 627–679 (2018).
- 552 **69.** Rouder, J. N., Morey, R. D., Speckman, P. L. & Province, J. M. Default bayes factors for ANOVA designs. *J. Math.*
553 *Psychol.* **56**, 356–374 (2012).
- 554 **70.** Rigoli, F. Masters of suspicion: A bayesian decision model of motivated political reasoning. *J. for Theory Soc. Behav.* **51**,
555 350–370 (2021).
- 556 **71.** Tajfel, H. E. *Differentiation between social groups: Studies in the social psychology of intergroup relations*. (Academic
557 Press, 1978).
- 558 **72.** Tajfel, H. E. & Turner, J. C. An integrative theory of group conflict. *The Soc. Psychol. Intergroup Relations* 7–24 (1986).
- 559 **73.** Kvam, P. D. & Pleskac, T. J. Strength and weight: The determinants of choice and confidence. *Cognition* **152**, 170–380,
560 DOI: 10.1016/j.cognition.2016.04.008 (2016).
- 561 **74.** Griffin, D. & Tversky, A. The weighing of evidence and the determinants of confidence. *Cogn. Psychol.* **24**, 411–435
562 (1992).
- 563 **75.** Tversky, A. & Kahneman, D. Belief in the law of small numbers. *Psychol. Bull.* **76**, 105 (1971).
- 564 **76.** Navarro, D. J., Perfors, A., Kary, A., Brown, S. D. & Donkin, C. When extremists win: Cultural transmission via iterated
565 learning when populations are heterogeneous. *Cogn. Sci.* **42**, 2108–2149 (2018).
- 566 **77.** Lynn, T., Muzellec, L., Caemmerer, B. & Turley, D. Social network sites: early adopters’ personality and influence. *J.*
567 *Prod. & Brand Manag.* (2017).
- 568 **78.** Stephen, A. T. & Lehmann, D. R. How word-of-mouth transmission encouragement affects consumers’ transmission
569 decisions, receiver selection, and diffusion speed. *Int. J. Res. Mark.* **33**, 755–766 (2016).
- 570 **79.** Mathew, B., Dutt, R., Goyal, P. & Mukherjee, A. Spread of hate speech in online social media. In *Proceedings of the 10th*
571 *ACM conference on web science*, 173–182 (2019).
- 572 **80.** Bastide, P. R., Deluca, L. S. & Do, L. M. User recommendations in a social media network (2018). US Patent 9,871,758.
- 573 **81.** Westfall, J., Van Boven, L., Chambers, J. R. & Judd, C. M. Perceiving Political Polarization in the United States: Party
574 Identity Strength and Attitude Extremity Exacerbate the Perceived Partisan Divide. *Perspectives on Psychol. Sci.* **10**,
575 145–158 (2015).
- 576 **82.** Thaler, R. H. & Sunstein, C. R. *Nudge: Improving decisions about health, wealth, and happiness* (Penguin, 2009).
- 577 **83.** Baujard, A., Gavrel, F., Igersheim, H., Laslier, J.-F. & Lebon, I. How voters use grade scales in evaluative voting. *Eur. J.*
578 *Polit. Econ.* **55**, 14–28 (2018).
- 579 **84.** Shiffrin, R. M. Memory search. *Model. human memory* 375–447 (1970).

Supplementary materials for “Rational inference strategies and the genesis of polarization and extremism”

Peter D. Kvam^{1,*}, Abhay Alaukik¹, Callie E. Mims², Arina Martemyanova¹, and Matthew Baldwin¹

¹University of Florida

²University of Southern Alabama

*Corresponding author: pkvam@ufl.edu

ABSTRACT

In this document, we provide a proof giving the expected distribution of information among decision makers and examine simulations based on manipulations of the decision making process. We apply our analyses of polarization and extremism to seven existing studies, showing that choice creates a sampling bias even in simple perceptual decisions. We additionally show that our results will hold for multi-alternative choice, such that polarization and extremism are not confined strictly to the domain of binary choice, and examine alternative model-based solutions to alleviating polarization and extremism.

1 Proof of rational polarization & extremism

A rational choice has to maximize expected utility in terms of both the goodness of a decision (accuracy, expected outcomes) as well as the time it takes to make a decision. The subjective utility that a decision maker experiences will balance both factors, and so a rational decision maker must identify the balance between desired accuracy and decision time that will yield the highest expected utility^{1,2}. To do so, a decision maker must compromise between spending more time on a choice and considering more information (lower utility from time, greater utility from accuracy) or spending less time and considering less information (greater utility from time, lower utility from accuracy). A decision maker should stop considering new information and make their choice when the expected benefit of considering one more piece of information (marginal improvement in accuracy) is outweighed by the loss of utility that accompanies spending the time required to consider that piece of information (marginal cost of time). Accuracy, quantified in terms of the odds of making a correct decision, is log-linearly related to the balance of support between options³. If each new piece of information can be expected to shift the balance of support toward one option (+1) or away from it toward another (-1), the marginal gain in accuracy (probability of correct) will decrease as the balance of support tips further and further toward one option.¹ A rational decision maker should therefore stop and make their choice when this marginal expected improvement in accuracy dips below the utility of the time it would take to gather that new piece of information. Put together, this decision process follows a random walk or diffusion process where the decision maker sets a threshold level of accuracy (balance of evidence) and gathers new information until that balance is reached⁵. The optimal balance of evidence may change over time⁶, but this accumulation-to-threshold strategy will maximize a decision maker's subjective expected utility across a wide range of choice environments.

Consider a choice scenario where a decision maker has two options they could select. They make their decision by gathering new pieces of information, which they use to sequentially update their beliefs about a stimulus. In order to achieve a particular level of accuracy while minimizing response times, the optimal strategy is to track the balance of support s for option A versus option B, and stop when this balance exceeds a threshold value θ . In the case where time is costly (spending time results in negative utility) or where there is a mix of difficulty in a sequence of choices, this threshold will decrease over the course of each decision⁶. Each new piece of information can increase or decrease the balance of support. Thus, the rational decision maker gathers new information piece by piece until either $s \geq \theta$ or $s \leq -\theta$ ^{5,7}. Here, we derive the balance of support at which a rational decision maker will stop and choose, showing that the resulting balance of support will follow a polarized distribution and create outlying extremists.

Suppose that new information is sampled from a true underlying distribution $f(x)$ that defines a distribution of evidence values it could take on $-\infty < x < \infty$ such that $\int_{-\infty}^{\infty} f(x) = 1$. The function $F(x)$ is the cumulative distribution function, determining the probability that the value of a new piece of information x_1 is less than or equal to x .

¹A similar claim can be said of multialternative choice, so our proof applies there as well⁴.

For now, we assume that all decision makers start in an unbiased state represented by the random variable $s(t)$, such that $s(0) = 0$.² Our goal is to derive the distribution of evidence states, specified by a probability density $g(x)$ over evidence levels x , that will result from a population of decision makers starting at evidence state $s_0 = 0$ and sampling new information from $f(x)$ until they reach state $\pm\theta$. We will designate the distribution of states of decision makers who finish at time t as $g_t(x)$, and the states of decision makers who have not finished before time t as a random variable $s(t)$.

The likelihood that a decision maker will stop after sampling a single piece of information x_1 , drawn randomly as $x_1 \sim f(x)$, is simply the probability of sampling a piece of information that moves the decision maker past one of the thresholds, $Pr(t = 1) = F(-\theta) + (1 - F(\theta))$. The distribution of evidence among those decision makers who have enough support to stop after this first piece of information, $g_1(x)$, will be

$$g_1(x) = f(x|x \leq -\theta) + f(x|x \geq \theta) \quad (1)$$

Meanwhile, those decision makers who do not stop after a single piece of evidence will be distributed between $-\theta$ and θ . The “unfinished” states at time point 1 will be s_1 , distributed according to a truncated distribution of $f(x)$:

$$Pr(s_1 = x) \sim \frac{f(x|-\theta < x < \theta)}{1 - Pr(t = 1)} \quad (2)$$

Those decision makers who have not made a selection yet will have to sample another piece of information from the source distribution $f(x)$, giving $s_1 + x_2$ where x_2 is a random variable drawn from $f(x)$. Some of these decision makers will then be able to stop and make a choice, if they have passed the threshold. Thus, the distribution of decision makers who have enough support to stop after two pieces of information will be

$$g_2(x) = Pr_x(x = s_1 + x_2) \cdot [Pr(x \leq -\theta) + Pr(x \geq \theta)] \quad (3)$$

where $x_2 \sim f(x)$. This procedure of adding a random variable from $f(x)$ to the existing distribution of states s_n , calculating the distribution below $-\theta$ / above θ , and truncating the remaining distribution of states can be carried out *ad infinitum* to get the distribution over evidence states that finish at each time point n .

$$g_n(x) = Pr_x(x = s_{n-1} + x_n) \cdot [Pr(x \leq -\theta) + Pr(x \geq \theta)] \quad (4)$$

Put together, this specifies a series of functions $g_1(x), g_2(x), \dots, g_n(x)$ that specify the distribution of evidence among decision makers that stop after the first, second, ..., n^{th} piece of information. Importantly, the probability of stopping on a particular step t , given that the participant has not stopped yet, is given by the integral of $g_n(x)$:

$$Pr(t = n | t > [n - 1]) = \int_{x=-\infty}^{\infty} g_n(x) dx. \quad (5)$$

Therefore, we can calculate the cumulative probability of stopping at or before step t as $G(t) = Pr(t = n | t > [n - 1]) \cdot (1 - G(t - 1))$, noting that $G(1) = \int g_1(x) dx$ and calculating subsequent values iteratively. Because $s(t)$ follows a Wiener process, there are well-established methods to determine $g(t)$ and $G(t)$ ^{8,9}. The overall distribution of evidence across decision makers at the conclusion of the decision process $S_t(x)$ can be then computed as

$$S_t(x) = \sum_{i=1}^{\infty} g_i(x) \cdot (1 - G(t)) \quad (6)$$

We could use $S_t(x)$ as the final distribution of evidence among decision makers who have finished deciding. However, it seems to make more sense for a decision maker’s opinion to be determined by the *average* of all of the pieces of information they have collected¹⁰, rather than as the *sum* of these pieces of information. (The summing rule makes the somewhat unreasonable prediction that there will be no decision makers whose beliefs lie between $-\theta$ and θ , and that participants’ views can be more extreme than any pieces of information they could consider.) The distribution of average evidence for decision makers arriving at a choice at time t is computed by simply mapping $g_t(x) \rightarrow g_t(\frac{x}{t})$, stretching or shrinking the x-scale by the number of pieces of information that had been drawn at time t .

²Variations on this assumption are explored in the model simulation section.

Polarization analyses

There were four main analyses carried out in the main text, whose results are presented in Table 1 and the Results. For the variance ratio analyses, we compared the variance of the observed information collected by participants against the expected variance of information that participants would have collected from a purely random sample (i.e., whose size did not depend on the stopping rule during choice). For Studies 1, 6, and 7, this was trivial: the exact stimulus values that could be generated on each trial were known, so we could use the objective variance on each trial along with the number of samples drawn on each trial to calculate the expected variance of the mean of the sample if it were representative. For these studies, the stimuli either generated multinomial random variables or normal random variables on each draw. Thus, we could calculate the expected variance of the mean of the multinomial / normal random variable for a given number of samples (given by the number the participant drew). This was calculated for every trial of every study, giving us a distribution of expected variance of the stimuli against which we could compare the observed variance of participants' samples.

As we might expect, the variance of observed and expected samples were generally quite close for these studies (Variance ratio in the first column of main text Table 1: Dots, DFE, and FGT).

Variance for the eye tracking studies was computed by taking the expected mean and variance for each individual fixation. It is the average of the squared differences between the expected value μ and the n attribute values a :

$$\sigma^2 = \frac{1}{n} \sum_i^n (a_i - \mu)^2$$

These were then used to calculate the expected value and variance of the mean for a trial, which was simply the sum of the the multinomial random variables describing the values we would expect them to sample on each fixation divided by the number of fixations. This allowed us to compute the expected variance of the mean for Studies 2-5 (Currency, Food, DFD, and Posters), against which we compared the observed average difference across fixations.

In contrast to the Dots, DFE, and FGT studies, the eye tracking studies tended to show particularly extreme patterns of variance. This seems to be due to a tendency toward autocorrelation in fixations – participants will look at the same option for multiple fixations (each typically around 250 ms), and tend to look at the option they are about to select for several fixations especially leading up to choice. This is referred to as the *gaze cascade*¹¹, where the entropy of fixations reduces dramatically, thus distorting the information they consider relative to an unbiased sampling distribution. This gaze pattern provides evidence for the use of the relative evidence (rational / optimal) strategies that underlie the pattern of polarization observed in the data¹².

Simulating Polarization

Although they validate our predictions related to both polarization and extremism, it is useful to generate predictions and characterize choices beyond the scope of empirical studies we present in the paper. We do so here with simulations, which permit us to explore the shape of distributions that manifest across many thousands of trials as a function of important parameters of the model.

The three most important parameters of a random walk / diffusion model related to distribution of evidence are the direction / rate of evidence accumulation (drift), the stopping rule (threshold), and the starting point bias (prior bias)¹³. We manipulate each of these parameters in our simulations in order to explore their effect on polarization and extremism. Drift describes the information that participants gather, both in terms of what the stimulus provides and in terms of what the participants seek out themselves. It denotes the strength of an average piece of information that a participant will sample during the information accumulation process. Bias in the drift rate, usually modeled as the addition of a fixed constant¹⁴, would reflect a decision maker's tendency to intentionally or unintentionally sample information that agrees with a particular position, including a confirmation bias in information seeking.

The second main parameter in our simulations, threshold, governs how much information a decision maker wishes to gather before they are certain enough of their position to make a choice. As we suggested in the main text, lower thresholds result in a decision maker basing their choices on fewer pieces of information, because they are willing to stop sooner. Classic manipulations of thresholds include incentives encouraging fast (low threshold) or accurate (high threshold) responses, such as time pressure or harsh punishments for incorrect answers^{2,15}. As we suggested in the main text, lower thresholds tend to result in a greater incidence of extremism. The reason for this will become apparent in the simulations.

The final parameter that we manipulate in our simulations is one that has not been discussed at length thus far. Start point bias describes the beliefs or preferences that a decision maker has before they gather any information at all, meaning that a decision maker needs some reason to favor one response over the other in absence of any new information. This could occur naturally in political spheres, where a new policy can have biases for or against it due to the nature of people's views on the issue or the way it is associated with political positions. However, start point bias tends to appear less frequently in experiments where all the stimuli are new and/or randomly generated. The main way to manipulate start point biases is to

incentivize responses on one side or another (for example, paying participants more if they respond “right” than “left”), or where participants get some predecision cue as to which option will be correct¹⁶.

Simulation results

The results of changing drift (accumulation bias), threshold, and start point (prior bias) are shown in Figure S1. We examine three main properties of the distributions: the proportion of responses favoring the right (+) or left (-) option, how widely spread decision makers are across the evidence space (width of the distributions) and the prevalence of extremists, which appear as “bumps” in the distributions far away from zero. Consider the top left panel, for example. There is an even distribution of responses favoring right and left alternatives (resulting from now drift bias), but responses are quite widely spread (high polarization) and there are substantial bumps in the distribution at around ± 0.5 (indicating a high prevalence of extremists).

Figure S1. Distribution of evidence among populations of decision makers with varying drifts (top vs bottom), thresholds (columns 1/3 vs 2/4), and starting point biases (left vs right). Low thresholds tend to promote both polarization and extremism, while start point biases have a large effect on extremism in particular.

The top row of Figure S1 corresponds to cases where there is no drift bias and the information that participants gather or seek out does not systematically favor one option or the other, while the bottom row shows cases where participants systematically consider more information for the response option on the right (positive drift). From the simulations, it is clear that changes in drift actually have relatively minimal effects on the shape of the positive or negative sides of the distribution; instead, it simply determines what proportion of responses favor the left or the right option (choice proportions are provided on either side of each panel). This indicates that extremism is probably not the result of information search biases, because manipulating drift does not inherently increase the proportion of decision makers who occupy extreme positions – the bumps do not increase in size relative to the overall density of their respective side of the evidence distribution.

The first and third columns of Figure S1 correspond to cases where decision makers adopt a low threshold for their decisions, while the second and fourth columns correspond to cases where decision makers adopt a high / strict threshold. This appears to have two effects: first, the overall distribution is more tightly distributed near zero, indicating a lower degree of polarization among decision makers. This should not be too surprising: the mean of a larger number of values (pieces of information) will have less variability than the mean of a small number of values (pieces of information), meaning that participants will wind up closer to the true average distribution of information in the environment as they increase their thresholds. The second effect of thresholds is to moderate the prevalence of extremists. As illustrated in the figure, there are larger bumps (more extremists) when thresholds are low than when they are high. This occurs because the first and second pieces of information that participants gather are less likely to terminate the decision making process, because they are unlikely to push a decision maker far enough in one direction to pass a high threshold. As a result, the probability of stopping with very little information ($Pr(t = 1) = F(-\theta) + (1 - F(\theta))$ in the proof above) is much lower when θ is large. Thus, one potential way to reduce the prevalence of both polarization and extremism is to induce higher thresholds in decision makers, through interventions encouraging accuracy and thorough consideration of evidence and discouraging fast and inconsiderate choice strategies. However, the overall shape of the distribution of information in the population will never truly reflect the input distribution. There is always a gap between those favoring the option on the left and the option on the right, meaning that polarization is an unavoidable consequence of decision making when agents implement an optimal balance-of-evidence stopping rule.

Finally, the left panels of Figure S1 correspond to situations with no starting point bias, while the right panels correspond to situations with a starting point bias toward the right-hand (+) response. As shown, starting point bias has a quite substantial effect on the prevalence of extremists – whereas they constitute a relatively small proportion of decision makers even in the no bias conditions, the proportion of extremists nearly triples when start point bias goes from 0 (left) to 0.4 (right). This occurs because decision makers who already favor positive responses require very little information to reach the positive boundary – many decision makers will stop after receiving a single piece of information favoring the option that they are already biased toward. As a result, the prior distribution of beliefs among decision makers will have a drastic effect on extremism – a wide start point distribution will result in many extremists on both sides, while a narrow start point distribution will reduce the prevalence of extremists on both sides.

In one sense, this is quite intuitive: when people start with more extreme beliefs already, they make up their minds very quickly and are likely to continue to hold extreme views. Our simulations additionally suggest that the decision making process will accentuate this extremism. Those that start with extreme views are likely to stop immediately after gathering information that makes them more extreme, and less likely to stop after sampling information that makes them less extreme. The decision making process therefore takes starting point biases and increases the bias in the population: the more decisions we make on related topics (where starting points could be expected to carry from one decision to the next), the more extreme our views become.

In many ways, this resembles interrogation and indoctrination strategies: have the subject admit to (decide in favor of) a small positive view, and then build on this small positive bias to make their views more extreme by iteratively offering them the chance to express (decide in favor of) more and more extreme views¹⁷. The similarities between indoctrination and polarization by sequential decisions are chilling, and our simulations suggest that they probably draw on a common mechanism: starting point bias. Further work on polarization and extremism should therefore consider the interaction between prior biases and decision making a central mechanism, worthy of deep consideration and exploration.

Verifying choice predictions with existing data

To examine whether polarization and extremism appeared in empirical work on decision making, we analyzed the results of seven different publicly-available decision making data sets. The tasks in these data sets varied in the type of stimuli and decision domain (preferential, inferential), but they had three important characteristics in common that made it possible to examine polarization and extremism: (1) the main goal of participants in the experiments was to make a series of decisions between binary choice options; (2) the true information provided by the stimulus was known, and tracked by the experimenters; and (3) each experiment recorded the final (average) set of information that participants had gathered on each trial. The reason behind the first criterion is that we are interested in the decision-making process and how it creates polarization. The other criteria are important because we need a benchmark against which to compare the empirical data. K-L divergence, and the plots shown in Figure S3, require an observed distribution (meaning the original experimenters had to record the information participants saw) and a reference / true distribution (meaning the stimulus distribution had to be known).

We should emphasize that this was not meant to be an exhaustive analysis of all studies that could have met our criteria, but a convenience sample of all the freely available studies on OSF we could find that included all the information we needed to carry out the analyses. A more detailed summary of the data from each study is provided below. We refer readers to the original studies / OSF pages for details on the samples, but we provide a short description of the seven experiments we analyzed here:

Study 1 - Dots¹⁸ In the dots study, participants were presented with a series of dots that appeared on the screen and asked to decide whether the dots were being pulled from a pool of dots that was 2/3 green and 1/3 pink or 2/3 pink and 1/3 green. From trial to trial, the proportion of dots that were actually pink or green (strength) was manipulated, as well as the rate at which these dots were presented (weight). Note that this presented a confound for threshold analyses, as thresholds in the original study were found to increase with higher values of strength and weight manipulations¹⁸, creating a violation of selective influence that would artificially increase the correlation between threshold estimates and the amount of evidence that participants gathered. This study included a total of 29 participants who each completed ~200 trials of the task.

Study 2 - Currency¹⁹ In this study, participants saw three currencies for each of two response options. Their task was to pick the set of three currencies that had the higher overall value (i.e., would yield the greatest total value sum). The currencies varied in type, including yen, pounds, and Q (made-up currency), and participants were given the exchange rates of each of these currencies to US dollars. Participants' gaze was tracked during the study so as to examine which of the outcomes (which of the six currencies) they were considering on each fixation. Positive and negative "sides" were randomly assigned such that the currencies on the top (bottom) conferred positive value while the currencies on the bottom (top) conferred negative value. This allowed us to examine the overall value of the pair of options: more positive indicated that the top (bottom) had greater value, while more negative indicated that the bottom (top) had greater value. On half the trials, the currencies displayed during a trial were all different (all three of yen, Q, and pounds for both options), and on the other half, the currencies displayed on a trial were all the same. We collapse across conditions for the analyses performed in this paper because they are not relevant to

our research questions. Participants responded by pressing one key (up) to select the first option on the screen, or another key (down) to select the second option. This study included a total of 46 participants, who each completed 84 trials of the task.

Study 3 - Food²⁰ In the food study, participants saw two snack foods displayed on either side of the screen and were instructed to decide which one they preferred. Before the decision trials, participants rated each of the 50 snack foods that could be displayed as stimuli in the experiments on a 1-9 scale in terms of desirability. These desirability values were used as inputs for the analyses we used in this paper – for example, a food with desirability value of 4 vs one with a desirability value of 6 would have a value difference of 2. As in the currency study, fixations were tracked and recorded. To get sample values from the eye tracking data, the desirability values of each option were centered so that fixating on the option with desirability 4 would have a value of -1, and fixating on the option with desirability of 6 would have a value of +1. A “representative” sample in this study was generated by randomly drawing the rated value of one snack or the other (50-50 chance). A total of 41 participants participated in this study, and each participant completed 100 trials.

Study 4 - Lotteries²¹ In the lotteries study, also referred to as a decisions from description study, participants were tasked with making decisions between two gambles. Each gamble had three equally likely outcomes, one each in the 10-30, 40-60, and 70-90 range. As in the currency study, values on the top / bottom were assigned positive / negative values so that positive outcomes for one option were contrasted against positive outcomes for the other option, providing a net expected value between choice options and allowing us to evaluate the variance of the values sampled on each fixation. A total of 54 participants completed 100 trials each of this study.

Study 5 - Posters¹⁹ The posters study was similar to the food study, except that Likert scales were gathered after the choice task rather than before. A “representative” sample in this study was generated by randomly drawing the rated value of one poster or the other (50-50 chance). A total of 53 participants completed the task and had usable data, and each one completed 100 trials of the study.

Study 6 - Decisions from experience²² In the decisions from experience study, participants were tasked with both learning about and choosing between two gambles. In contrast to the decisions from description study, participants had to learn about each of the gambles by sampling from each one, rather than viewing the probabilities directly. For example, they could experience a series of samples [\$4, \$0, \$4, \$4, \$0, \$4] from a gamble that was 75% chance of \$4 and 25% chance of \$0. The different choice problems varied in terms of the number, magnitude, and probability of the different outcomes that participants could receive from either choice option. Once participants indicated they were finished sampling, they chose one option or the other. The exact sequence of draws obtained by a participant in this task was recorded and used for our analyses – as in the other studies, one option was coded as negative and the other as positive in order to determine the expected and observed variance / mean of the samples that participants drew. A “representative” sample in this study was generated by drawing randomly between the available decks (50-50 chance) and generating outcomes according to their objective probabilities. A total of 104 participants completed the study, and each one completed 228 trials.

Study 7 - Flash gambling²³ In the flash gambling study, participants were tasked with determining which of two flashing patches of dots had a greater mean number of dots (inference) or which one they would prefer to take a random draw from, given that more dots meant a greater payoff (preference). One patch of dots had a constant number of dots (130), and the dots simply changed positions from frame to frame. For the other patch of dots, a random number was generated with a known mean (e.g., 140) and standard deviation (e.g., 20) on each frame, and the resulting number of dots was displayed. The difference between the certain and uncertain patches therefore followed a normal distribution, which was used to define a “representative” draw for our analyses. However, the patches were refreshed at 60 Hz, making it very difficult to stop the stimulus immediately after obtaining a streak of information favoring one option or another. As a result, this study showed rather weak polarization and extremism results (Tables S1 & S2). A total of 113 participants each completed 960 trials of the experiment.

Each of the existing data sets we analyzed came from an open data set hosted on the Open Science Framework (OSF). The dots data are provided at osf.io/ba5c7/; the currency, food, poster, and lottery / DFD data are provided at osf.io/mvk95/²⁴; the decisions from experience / DFE data are provided at osf.io/ngc45/; and the flash gambling task data (Study 1) are provided at osf.io/g7a49/. These studies were chosen so that we knew the data were (1) coming from an incentivized decision making task, (2) tracked the information that participants considered during the decision process, and (3) had all of the necessary information to compute the expected mean and variance of evidence of a random sample for the stimuli on each trial. We summarize the essential characteristics of the methods from each study here, but more detailed descriptions and explanations can be found in the original papers.

Studies 1, 6, and 7 were experiments in which participants received information piece by piece, allowing us to directly compare the information they saw against the information we would expect based on a random draw from the stimulus. Studies 2-5 were eye tracking studies that compared the value of stimulus attributes that participants looked at (observed) against the expected value of the stimuli that participants would have sampled if they looked randomly at stimulus attributes for the same number of fixations. These three characteristics allowed us to examine how the true information generated by the stimulus compared to the information that participants had gathered at the end of the decision making process.

All of the analyses presented in this paper used a Bayesian method for estimation or hypothesis testing (model comparison) with uninformative priors on all parameters (wide normal for unrestricted range, gamma for positive-only, or uniform for range-restricted parameters) that allowed the data to supplant the priors in determining the results. Analyses were implemented in JAGS²⁵, using MCMC sampling to estimate the posterior distribution of the statistic of interest (variance ratio, correlation, etc.). For all analyses, we present the mean estimate and the 95% highest density interval (HDI) for each effect²⁶.

To test for polarization, we compared the true stimulus information (representative samples) against the information that participants actually collected in each of the data sets. The first analysis we carried out directly compared the distributions of observed samples collected by participants against the distribution of samples we would expect to obtain from a random, representative sample. This allowed us to say how much participants differed from the – typically unimodal – distributions of evidence that a true random sample would provide.

To test for extremism, we compared the variance of the samples that participants gathered against the variance of a representative sample. This provided us with a measure of how extreme participants' sampled information was, as a ratio between the views they formed and the true information in their environment. We present a formal description of each of these analyses next.

Polarization analyses

The analyses for polarization followed those presented in the main text, using the Kullback-Leibler divergence [KLD]²⁷ to measure the difference between the expected unimodal distribution of evidence and the evidence that participants actually collected during the experiment. The K-L divergence was approximated by calculating the probability density of the expected and observed distributions of evidence for a large number of points – 1001 evenly-spaced points from the minimum to the maximum of the observed data points or of the possible outcomes (for bounded scales like in the Dots study). This was used to approximate the integral

$$KLD[p(y)||p(x)] = \int_i^N p(\mathbf{y}_i) \log \left(\frac{p(\mathbf{y}_i)}{p(\mathbf{x}_i)} \right)$$

where y is the observed distribution and x is the expected distribution, with their probability densities at location i given by $p(x_i)$ and $p(y_i)$.

The density of the observed distribution was calculated by passing an optimal kernel density estimator over the observed data (the default values of kernel bandwidth in MATLAB's `ksdensity` being optimal for normal distributions²⁸) to calculate a probability density at each of the points i .

However, this only gives a point estimate of the degree of polarization in the data. To construct a 95% credible interval around this estimate, we generated 10,000 simulated data sets from the expected distributions. This was carried out by first computing the true probability density function for the expected distribution (needed for the KLD calculation anyway), then using the cumulative probability density to perform inverse transform sampling²⁹ to sample N random draws from the representative / expected distribution, where N was the number of data points in the observed data.

For each of the 10,000 simulated data sets, which matched the observed data in size, we then calculate the KLD between the simulated data and the true expected sampling distribution using the same kernel density approximation method. This gave us a sample of 10,000 KLD values that described the distribution of KLDs we would obtain if the data were generated from the expected distribution. Once this was computed, we could compare the KLD values between the observed data and the 10,000 simulated data sets. This allowed us to construct the mean difference and a 95% credible interval around the difference between the observed KLD and the expected KLD (from the simulated data) from a representative distribution. The results of this analysis for each study are reported in Table S1 below.

Extremist analyses

To examine the degree of extremeness in participants' views, we compared the expected variance of the stimulus information (if a participant drew n samples at random, what would we expect the variance to be?) against the true variance of the information collected by participants (what was the variance of the n samples they actually collected?). We refer to this as the *variance ratio* analysis, specifying the ratio of observed variance to expected variance. The variance ratio for each experiment is provided in the first column of Table S2: Variance ratios greater than 1 indicate that participants collected more extreme information/were more polarized than we would expect purely by chance, while variance ratios less than 1 indicate that a participant collected less extreme information/were less polarized than expected.

In addition to the variance ratio analyses, we also calculated the extremeness of each trial and examined how individual differences in extremism were related to other patterns of behavior on the task. Specifically, we looked at correlations between the amount information participants gathered, the extremeness of the information they collected on each trial, and the rate

Figure S2. Relationship between sample size (x) and the average Bayes factor in favor of the null hypothesis (y). The average across 10,000 simulates is shown as the dark blue line, while the shaded region indicates the 20th and 80th percentiles of Bayes factors for each simulated sample size.

at which they made their decisions (response time). The degree of extremeness of the evidence participants gathered on a particular trial was defined as the absolute value of the average evidence they collected on each trial minus the true mean of the stimulus on that trial. This extremeness was correlated with the number of samples gathered on each trial (Information, Column 2 of Table S2) and the response time on each trial (Speed, Column 3 of Table S2).

Each of these correlation analyses were performed using Bayesian methods. The values for amount of information and response times were standardized so that the resulting correlations would be on a $[-1, 1]$ scale. The prior for the correlation between variables was a uniform distribution $U(-1, 1)$, and they were estimated using JAGS²⁵. The posterior was generated using 4 chains of 5000 samples, with 500 burn-in steps in each chain and starting points generated randomly from the prior. These same settings were used to estimate the polarization effects as well, except that the variance ratio had a wide uniform prior $U(0, 1000)$ and the extremity had a wide normal prior $N(0, 1000)$. This allowed us to approximate the posterior distribution of each of the effects – the means and 95% Highest Density Intervals (95% HDIs) of the effects are reported in Table S2.

Sample size rationale

The existing data sets had already been collected, and therefore there is not much to be done about their size or composition. Despite this, they are quite large and thorough: put together, there were a total of 440 participants and over 150,000 trials across these data sets. They spanned inferential and preferential choice, eye tracking and sampling paradigms for gathering information, and both description and experience-based types of information. This wide-ranging set of studies suggests that the phenomena we observe are pervasive, and in particular that polarization and the presence of extremists are consequences of adopting relative-evidence decision strategies regardless of the particular domain of choice.

For the new studies presented in the main text, we aimed to gather sufficient data to be able to relate the individual difference measures we collected (Dogmatism, Personal Need for Closure) to individual differences in sampling behavior such as the number of clicks of degree of polarization. We determined that we would like to be able to conclude in favor of the null hypothesis (no relationship between measures) at least 80% of the time when the null hypothesis was in fact true, an inference that is made possible by the use of Bayesian statistics³⁰. To estimate how many participants would be needed, we created 10,000 artificial data sets at differing sample sizes, shown in Figure S2, where the relationship between two variables was zero (i.e., the null hypothesis was true). For each artificial data set, we computed the Savage-Dickey Bayes factor³¹ quantifying the degree of support for or against the null hypothesis. Bayes factors greater than 3 (less than $1/3$) indicate solid support for (against) the null. We determined that 85 participants was sufficient to ensure that we would support the null with a Bayes factor of 3 or more in at least 80% of our analyses.

In order to obtain a more representative sample of U.S. participants (which we focused on because of the nature of the political topics) than we might expect from a university undergraduate sample, we recruited online through Prolific Academic.

Table S1

Study	Expected KLD	Observed KLD	Difference	95% HDI
Dots	3.85	4.09	0.25	[-0.75, 1.27]
Currency	0.02	0.74	0.73	[0.72, 0.73]
Food	0.03	0.31	0.28	[0.27, 0.29]
DFD	0.0069	0.4	0.39	[0.38, 0.39]
Posters	0.028	1.28	1.25	[1.24, 1.26]
DFE	0.018	2.56	2.54	[2.53, 2.54]
FGT	0.0004	0.0086	0.0082	[0.0081, 0.0084]

Table S2. Table of mean effects [95% HDIs] related to extremism for each of the seven studies. The variance ratio is the relative variance of the observed versus expected samples (greater than 1 indicates polarization). For extremists, we examined the relationship between the extremeness of evidence on a particular trial and the amount of information gathered on that trial (Information) and the response times on those trials (Speed).

Study	Extremism	Relation to other measures	
	Variance ratio	Information	Speed
Dots	1.21 [1.16, 1.26]	-.39 [-.42, -.36]	-.41 [-.44, -.39]
Currency	7.24 [6.86, 7.69]	-.11 [-.15, -.07]	-.07 [-.11, -.03]
Food	3.21 [3.08, 3.36]	-.04 [-.07, -.01]	-.14 [-.17, -.10]
DFD	3.28 [3.09, 3.51]	-.10 [-.15, -.05]	-.14 [-.19, -.10]
Posters	8.51 [8.09, 9.98]	-.04 [-.08, -.00]	-.17 [-.20, -.13]
DFE	2.13 [2.10, 2.17]	-.19 [-.20, -.17]	-.23 [-.24, -.21]
FGT	0.78 [0.77, 0.79]	-.25 [-.26, -.24]	-.55 [-.56, -.55]

Results

The analysis for polarization compared the KLD of the observed data relative to the true probability density (observed KLD) against the results of 10,000 matched- N random samples from the true distribution relative to the true probability density (expected KLD). This allowed us to estimate a 95% highest density interval on the differences between the observed polarization and the degree of polarization we would expect to obtain just by chance. The results of this analysis are shown in Table S1 for each study we re-analyzed.

Out of the seven studies, six showed a difference in KLD between the observed data and the random samples from the true distribution. For the Currency, Food, DFD, Posters, and DFE studies, the KLD was several times greater than we would expect purely by chance, resulting in stark differences between the observed KLD and the distribution of expected KLDs. This suggests a large degree of polarization in the information collected by participants in these studies, who gathered information in a way that biased them away from the center of the distribution / more moderate information. For the Dots study, the observed KLD was larger than the expected KLD, but the random samples from the true distribution had sufficient variability that the observed KLD was not large enough to constitute a credible difference. The FGT study also showed a credible KLD difference, but this may actually run in the opposite direction to the polarization effect observed in the other studies (see Figure S3) and indicate that the difference between distributions (which was quite small anyway) was driven by more observed samples near 0 than would be expected from a representative sample.

In addition to the polarization analyses, we examined the extremeness of the information that participants collected in each study, defined as the average value of all the pieces of information they collected during a trial. Aggregate estimates of extremeness were created by combining these trial-level outcomes into individual differences (by aggregating within participants) and into an overall extremeness score for each of the experiments (by aggregating across participants). The variance of the study-aggregated samples was compared to the expected variance from a representative sample, giving the first *variance ratio* column of Table S2. The individual-level extremism measures (aggregated only within subjects) were then correlated with important outcomes like how many pieces of information each person gathered (Information) and how quickly they made their choices (Speed). These linear correlations are presented in the second and third columns of Table S2, respectively.

In six of the seven studies, we found that the variance in the information gathered by participants was greater than would be expected given a representative sample from the distribution (Figure S3). The only study that showed a pattern of results that was inconsistent with a pattern of polarization was the Flash Gambling Task (FGT) study. Because information arrived rapidly in this study (the screen updated at 20 Hz), participants were likely unable to consider every single piece of information.

Figure S3. Distribution of observed information gathered by participants (blue histograms) compared to the expected distribution of information that a participant *should* have, given representative draws from the true distributions.

Furthermore, even if participants had made up their minds after gathering an extreme piece of information, information arrived continuously until they made their response. If the intervening period were around 200-300 ms³², then they would have gathered at least 4-6 additional samples of evidence, diluting the extremity of the pool of information they appear to have considered before the end of the trial. This would naturally wash out any effects of the participants' stopping rules and negate the typical extremism effect that we observed in other studies.

The remaining studies showed a robust and large extremism effect, with the estimated variance of the samples participants collected ranging from 121% to 851% of the variance we would expect from random sampling. To illustrate this effect visually, the patterns of polarization in these studies are also shown in Figure S3. The expected distribution of information that we would anticipate from a representative sample is shown in orange, while the observed distribution of information collected by participants is shown as a blue histogram. Clearly, the blue histograms show a more widely dispersed pattern of information than the expected distributions (orange lines).

Put together with the polarization (Table S1) and extremism (Table S2) analyses, it is clear that the decision strategies participants implemented in these studies – with the possible exception of the FGT study – created a bias in the information they collected. These results constitute strong evidence for the effect of decision goals on polarization, suggesting that polarization and extremism can arise without motivated bias. Instead, polarization and extremism arise as a function of the information decision makers collect; and crucially, occurs even on simple perceptual tasks.

Extremists in the distribution

Although the presence of polarization is an important implication of widespread use of optimal choice strategies, it is not the only interesting phenomenon that arises as a consequence of implementing these strategies to make decisions. Among populations of decision makers, there will naturally be variation in the value that people place on making the correct (accurate, informed) decision versus making a timely decision. This trade-off between speed and accuracy is reflected in the setting of the threshold in dynamic models of decision making^{2,15}. As illustrated in the proof and simulations, some decision makers executing optimal strategies favoring speed will appear as “lumps” at the poles of the information distribution. These individuals tend to be decision makers who make up their minds after seeing only a few pieces of extreme information, emphasizing speedy sampling over thorough sampling when making decisions. They tend to be both polarized (have no views near the center of the distribution, high individual-level KLD) and adopt extreme views (high variance in collected information). For simplicity, however, we refer to them as *extremists*, as they tend to appear at the tails of the distributions of information and possess distorted views of the information that is in their environments. Extremism appears when a decision maker who favors speedy decisions (low thresholds) encounters a small set or even a single piece of extreme information. This piece of information allows them to immediately cross their threshold and favor $A \succ B$ or $B \succ A$. They can therefore stop gathering new information, meaning that the extreme view provided by the lone piece of extreme information is never combined or tempered with conflicting evidence.

Evaluating our first measure of extremism, we found the averages of smaller samples tended to be more volatile, and thus the extremists were likely to have collected less information than their peers (summarized in the Information column in Table S2). In all seven studies, this is evidenced by a credible (95% HDI excluding zero; significant, in classical statistical terms) negative correlation between the number of pieces of information and the overall extremity of that information. Similarly, our second measure found that in all seven studies, there was a negative relationship between response time and information

extremity, suggesting that participants with extreme views tended to make faster decisions during the experiments (Speed column in Table S2).

Another way to understand extremism is to explore underlying individual differences in the cognitive mechanisms supporting choice, which might offer insights into what made these particular individuals become extremists. We suggest that low thresholds, indicating a preference for speedy as opposed to well-informed decisions, is more common among extremists (Figure S1). To investigate this, we applied a simple diffusion model¹³ to predict choices and response times in terms of the strength of information favoring one option over another (drift), priority for speed or accuracy (threshold), and the length of time spent on processes unrelated to choice such as encoding stimuli and making a response (non-decision time). For each trial of each study, the drift rate was fixed based on the true average difference in value between stimuli. Threshold and non-decision time were permitted to vary across participants (none of the studies manipulated speed / accuracy instructions), so that each person was assigned a unique value for their threshold. This allowed us to compare the average extremeness of an individual participant's responses (extremism) against a rough estimate of the threshold they used to make choices. This analysis was carried out for the five studies it was possible to apply the diffusion model³.

This disentangled performance for each participant on each task into the strength of information provided by the stimulus (drift), the amount of time on each trial devoted to non-decision processes (non-decision time), and the amount of information participants wanted to collect before deciding (threshold). Because the first two measures are primarily related to the stimulus and to participants' motor abilities, we looked only at the threshold parameter. In four of the five studies, our analysis of individual differences in decision thresholds showed that there was a credible (significant) negative correlation between extremism and thresholds estimated for a diffusion model: Food ($b = -.04$, 95% HDI = $[-.05, -.03]$), DFD ($b = -.08$, 95% HDI = $[-.11, -.06]$), Posters ($b = -.06$, 95% HDI = $[-.08, -.04]$), and FGT ($b = -.04$, 95% HDI = $[-.05, -.03]$). While not overwhelming evidence, these results do suggest that decision makers with more extreme views tended to favor fast decisions over well-informed ones. We explore potential individual differences that may underlie this extremism effect in the new study, described in the following sections.

Alternative solutions to polarization

Although the estimation manipulation is the main approach we used to try to reduce polarization and extremism in information sampling, there are other possible approaches that may or may not work. One seemingly obvious strategy to reduce polarization in binary choices like those we present here, would be to simply offer more choices. The authors ourselves were hopeful that multi-alternative choice would curb polarization and extremism, thinking that polarization was an issue of having two opposing options and that adding a third could eliminate or reduce it. If true, the addition of another viable alternative to reduce polarization could provide justification for eliminating two-party systems (resulting from first-past-the-post) or for ranked-choice voting protocols as interventions.

Unfortunately, this was not the case – a similar pattern of polarization and extremism occurs beyond binary choice as well. Optimal choice strategies for these situations are provided in work by^{4,33,34}: essentially, a decision maker's beliefs in optimal trinary choice can be described by a location on a triangle, and by a multidimensional simplex when there are 4+ options⁴. To make a decision, the decision maker starts on the interior of the triangle (simplex) and gathers information, changing their beliefs accordingly until they cross one of the edges of the triangle (sides of a simplex). The size of the triangle (simplex) is determined by the height of the thresholds. This means that extreme information is still more likely than moderate information to terminate the decision process, with the end result being the same as in binary choice.

This is illustrated in Figure S4: a fully symmetric distribution of input choice evidence creates three “pockets” of posterior evidence favoring each option, centered around those favoring A, those favoring B, and those favoring C. There are very few decision makers between A and B, B and C, and A and C, and yet fewer who have a near-even view of all three. In fact, across more than 100,000 simulations shown in Figure S4, we had no instances of decision makers who had a balanced or even nearly-balanced view of the three options (white area of the figure). More choice does not mean less polarization.

Absolute evidence

One final approach that may help reduce polarization and extremism is for people to represent support for different choice options separately, rather than as a balance between the options they have available. In accumulator models of decision making^{35,36}, people are thought to implement an *absolute stopping rule*, where they gather support for each of the available choice options until the degree of support for one of the alternatives reaches a threshold. Despite making nearly identical predictions for the shape of response time distributions³⁷, accumulator models differ substantially in the pattern of information accumulation we observe. A simulation of these strategies is shown in Figure S5.

³We could not include the Dots study, as it contained a manipulation that artificially induced a relationship between the amount of information and speed of decision making¹⁸, nor the DFE study, since it has a discrete sampling structure which meant that response times were not commensurate with the number of pieces of information drawn.

Figure S4. Distribution of beliefs (evidence) among a population of decision makers who have made a decision among three options A, B, and C.

Figure S5. Simulated data from an accumulator model with no interaction between accumulators (left) and with lateral inhibition between them (right), compared against the distribution of information we would expect from a relative stopping rule (dotted black line).

As shown, an absolute stopping rule does have the potential to reduce polarization similar to estimation goals. By separating support into support for A, support for B, etc., information does not have to be extreme in order to trigger a choice. Instead, the support for each option is the sum of any small, moderate, or large increase in support; a decision maker does not have to be sensitive to the balance of support between options, just the overall degree of support for each one. As a result, extreme information loses some of its value in triggering a decision and the overall distribution of information that (simulated) participants collect reverts to a more representative one as shown in the left panel of Figure S5.

Whether representing support for different options separately will reduce polarization, however, depends on whether there are any interactions between support for different options. In many accumulator models of decision making, support for one option decreases the rate at which support for another option can increase, referred to as *lateral inhibition*^{38,39}. These interactions between options can actually exacerbate the problem of polarization: if a person already strongly supports one option, they will not gather support for contrary options. As a result, people get “funneled” into strong support for one option or another as soon as they gather enough one-sided support for an option. As a result, they wind up in a tri-modal distribution of support: those who favor option A and will not consider information for option B, those who favor option B and do not consider information for option A, and those who maintain a relatively even balance of support for A and B throughout the accumulation process.

The result of many simulations of this strategy is shown in the right panel of Figure S5: instead of removing polarization, absolute stopping rules with interacting accumulators can exacerbate the problem of extremism. The effect of absolute stopping rules and the moderating effect of lateral inhibition are certainly worthy of investigation in future interventions seeking to reduce polarization and extremism.

Fixed sampling

Another way to reduce polarization is to simply force decision makers to take some number n pieces of information before stopping, as in fixed-sampling paradigms^{40,41}. If the information that decision makers gather is independent and identically distributed (i.i.d.) across samples, as is typically assumed, decision makers will obtain a representative sample of the true information in the environment and thus polarization will not occur. However, this approach is paternalistic and thus potentially very difficult to implement^{42,43}, as it takes the autonomy of information collection away from decision makers⁴⁴. In future studies, it would be interesting to test more benign versions of such interventions such as providing a large initial sample of information for decision makers to sort through before gathering additional samples.

References

1. Savage, L. J. *The Foundations of Statistics* (John Wiley & Sons, Inc, New York, NY, 1954).
2. Wickelgren, W. A. Speed-accuracy tradeoff and information processing dynamics. *Acta Psychol.* **41**, 67–85, DOI: 10.1016/0001-6918(77)90012-9 (1977).
3. Wald, A. & Wolfowitz, J. Bayes solutions of sequential decision problems. *Proc. Natl. Acad. Sci. United States Am.* **35**, 99–102 (1949).
4. Kvam, P. D. A geometric framework for modeling dynamic decisions among arbitrarily many alternatives. *J. Math. Psychol.* **91**, 14–37 (2019).
5. Bogacz, R., Brown, E., Moehlis, J., Holmes, P. & Cohen, J. D. The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* **113**, 700–765, DOI: 10.1037/0033-295X.113.4.700 (2006).
6. Malhotra, G., Leslie, D. S., Ludwig, C. J. & Bogacz, R. Time-varying decision boundaries: insights from optimality analysis. *Psychon. bulletin & review* **25**, 971–996 (2018).
7. Edwards, W. Optimal strategies for seeking information: Models for statistics, choice reaction times, and human information processing. *J. Math. Psychol.* **2**, 312–329, DOI: 10.1016/0022-2496(65)90007-6 (1965).
8. Diederich, A. & Busemeyer, J. R. Simple matrix methods for analyzing diffusion models of choice probability, choice response time, and simple response time. *J. Math. Psychol.* **47**, 304–322, DOI: 10.1016/S0022-2496(03)00003-8 (2003).
9. Ratcliff, R. A theory of memory retrieval. *Psychol. Rev.* **85**, 59–108, DOI: 10.1037/0033-295X.85.2.59 (1978).
10. Turner, B. M., Gao, J., Koenig, S., Palffy, D. & McClelland, J. L. The dynamics of multimodal integration: The averaging diffusion model. *Psychon. Bull. & Rev.* **24**, 1819–1843 (2017).
11. Shimojo, S., Simion, C., Shimojo, E. & Scheier, C. Gaze bias both reflects and influences preference. *Nat. neuroscience* **6**, 1317–1322 (2003).

12. Mullett, T. L. & Stewart, N. Implications of visual attention phenomena for models of preferential choice. *Decision* (2016).
13. Ratcliff, R., Smith, P. L., Brown, S. D. & McKoon, G. Diffusion decision model: Current issues and history. *Trends Cogn. Sci.* **20**, 260–281 (2016).
14. Diederich, A. & Busemeyer, J. R. Modeling the effects of payoff on response bias in a perceptual discrimination task: Bound-change, drift-rate-change, or two-stage-processing hypothesis. *Percept. & Psychophys.* **68**, 194–207, DOI: 10.3758/BF03193669 (2006).
15. Heitz, R. P. The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Front. Neurosci.* **8**, 150 (2014).
16. Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W. & Forstmann, B. U. Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *J. Neurosci.* **32**, 2335–2343 (2012).
17. Hinkle, L. E. & Wolff, H. G. Communist interrogation and indoctrination of enemies of the states: Analysis of methods used by the communist state police (a special report). *AMA Arch. Neurol. & Psychiatry* **76**, 115–174 (1956).
18. Kvam, P. D. & Pleskac, T. J. Strength and weight: The determinants of choice and confidence. *Cognition* **152**, 170–380, DOI: 10.1016/j.cognition.2016.04.008 (2016).
19. Edmunds, C. E., Bose, D., Camerer, C. F., Mullett, T. L. & Stewart, N. Experimental investigations of attention interaction in preferential choice. <https://osf.io/mvk95/> (2020).
20. Bhatia, S. & Mullett, T. L. Similarity and decision time in preferential choice. *Q. J. Exp. Psychol.* **71**, 1276–1280 (2018).
21. Mullett, T. L., Ravichandar, N. & Stewart, N. Attention and accumulation in risky choice. <https://osf.io/mvk95/> (in prep).
22. Kellen, D., Pachur, T. & Hertwig, R. How (in) variant are subjective representations of described and experienced risk and rewards? *Cognition* **157**, 126–138 (2016).
23. Pleskac, T. J., Yu, S., Hopwood, C. & Liu, T. Mechanisms of deliberation during preferential choice: Perspectives from computational modeling and individual differences. *Decision* **6**, 77–107 (2019).
24. Edmunds, C. E., Bose, D., Camerer, C. F., Mullett, T. L. & Stewart, N. Accumulation is late and brief in preferential choice. *PsyArXiv* DOI: 10.31234/osf.io/sa4zr (2020).
25. Plummer, M. JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*, vol. 124, 10 (Vienna, Austria., 2003).
26. Kruschke, J. K. *Doing Bayesian Data Analysis: A Tutorial with R, JAGS, and STAN* (Academic Press, 2014).
27. Kullback, S. & Leibler, R. A. On information and sufficiency. *The annals mathematical statistics* **22**, 79–86 (1951).
28. Bowman, A. W. & Azzalini, A. *Applied smoothing techniques for data analysis: the kernel approach with S-Plus illustrations*, vol. 18 (OUP Oxford, 1997).
29. Vogel, C. R. *Computational methods for inverse problems* (SIAM, 2002).
30. Lee, M. D. & Wagenmakers, E.-J. *Bayesian Cognitive Modeling: A Practical Course* (Cambridge university press, 2014).
31. Wagenmakers, E. J., Lodewyckx, T., Kuriyal, H. & Grasman, R. Bayesian hypothesis testing for psychologists: A tutorial on the Savage-Dickey method. *Cogn. Psychol.* **60**, 158–189, DOI: 10.1016/j.cogpsych.2009.12.001 (2010).
32. Ratcliff, R. & McKoon, G. The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Comput.* **20**, 873–922, DOI: 10.1162/neco.2008.12-06-420 (2008).
33. Link, S. W. & Heath, R. A. A sequential theory of psychological discrimination. *Psychometrika* **40**, 77–105, DOI: 10.1007/BF02291481 (1975).
34. Tajima, S., Drugowitsch, J., Patel, N. & Pouget, A. Optimal policy for multi-alternative decisions. *Nat. neuroscience* **22**, 1503–1511 (2019).
35. Smith, P. L. & Vickers, D. The accumulator model of two-choice discrimination. *J. Math. Psychol.* **32**, 135–168, DOI: 10.1016/0022-2496(88)90043-0 (1988).
36. Brown, S. D. & Heathcote, A. The simplest complete model of choice response time: Linear ballistic accumulation. *Cogn. Psychol.* **57**, 153–178, DOI: 10.1016/j.cogpsych.2007.12.002 (2008).
37. Donkin, C., Brown, S., Heathcote, A. & Wagenmakers, E.-J. Diffusion versus linear ballistic accumulation: Different models but the same conclusions about psychological processes? *Psychon. Bull. & Rev.* **18**, 61–69, DOI: 10.3758/s13423-010-0022-4 (2011).

- 545 **38.** Usher, M. & McClelland, J. L. The time course of perceptual choice: The leaky, competing accumulator model. *Psychol.*
546 *Rev.* **108**, 550–592, DOI: 10.1037/0033-295X.108.3.550 (2001).
- 547 **39.** Usher, M. & McClelland, J. L. Loss aversion and inhibition in dynamical models of multialternative choice. *Psychol. Rev.*
548 **111**, 757–769, DOI: 10.1037/0033-295X.111.3.757 (2004).
- 549 **40.** Busemeyer, J. R. Decision making under uncertainty: a comparison of simple scalability, fixed-sample, and sequential-
550 sampling models. *J. Exp. Psychol. Learn. Mem. Cogn.* **11**, 538 (1985).
- 551 **41.** Fried, L. S. & Peterson, C. R. Information seeking: Optional versus fixed stopping. *J. Exp. Psychol.* **80**, 525 (1969).
- 552 **42.** Grüne-Yanoff, T. & Hertwig, R. Nudge versus boost: How coherent are policy and theory? *Minds Mach.* **26**, 149–183
553 (2016).
- 554 **43.** Hertwig, R. & Ryall, M. D. Nudge versus boost: Agency dynamics under libertarian paternalism. *The Econ. J.* **130**,
555 1384–1415 (2020).
- 556 **44.** Deci, E. L. & Ryan, R. M. Human autonomy. In *Efficacy, agency, and self-esteem*, 31–49 (Springer, 1995).