

The Politically Motivated Reasoning Paradigm, Part 1: What Politically Motivated Reasoning Is and How to Measure It

DAN M. KAHAN

Abstract

Recent research identifies *politically motivated reasoning* as the source of persistent public conflict over policy-relevant facts. This essay, the first in a two-part set, presents a basic conceptual model—the *Politically Motivated Reasoning Paradigm*—and an experimental setup—the *PMRP design*—geared to distinguishing the influence of PMRP from a truth-seeking Bayesian process of information processing and from recurring biases understood to be inimical to the same. It also discusses alternative schemes for operationalizing “motivating” group predispositions and the characteristics of valid study samples for examining this phenomenon.

THE NEW POLITICS OF “FACT POLARIZATION”

Polarization over questions of fact is a signature feature of contemporary political life. Citizens divided on the relative weight of “liberty” and “equality” disagree less intensely on the justice of progressive taxation (Moore, 2015) than on the reality of human-caused global warming (Frankovic, 2015). Democrats and Republicans argue less strenuously about whether public schools should permit “voluntary prayer” (Smith, Marsden, & Hout, 2014) than about whether permitting citizens to carry concealed handguns increases or decreases homicide rates (Newport, 2015).

These are admittedly complex questions. However, they are empirical ones. Values cannot supply the answers; only *evidence* can.

Whether humans are heating the earth and concealed-carry laws increase crime, moreover, turn on wholly distinct bodies of evidence. There is no logical reason for positions on these two empirical issues—not to mention myriad others, including the safety of underground nuclear-waste disposal,

the deterrent impact of the death penalty, the efficacy of invasive forms of surveillance to combat terrorism—to cluster *at all*, much less form packages of beliefs that so strongly unite citizens of one set of outlooks and divide those of opposing ones.

However, there is a *psychological* explanation. Or at least a very strong candidate, the emergence of which has supplied an energizing focus for decision science research.

That explanation is *politically motivated reasoning* (Jost, Hennes, & Lavine, 2013). Where positions on some policy-relevant fact have assumed widespread recognition as a badge of membership within identity-defining affinity groups, individuals can be expected to selectively credit all manner of information in patterns consistent with their respective groups' positions. The beliefs generated by this form of reasoning excite behavior that expresses individuals' group identities. Such behavior protects their connection to others with whom they share communal ties (Sherman & Cohen, 2006).

Indeed, what an ordinary citizen believes about the effect of private gun possession, the contribution of humans to climate change, and like facts will typically have no meaningful impact on the risks these states of affairs pose or on adoption of policies relating to them. The reliable activation of affective stances that convey group allegiance will be the *only* use most citizens have for such beliefs. In such circumstances, politically motivated reasoning can be understood to be perfectly rational (Kahan, in press).

This essay, the first of the two, will synthesize the research supporting this account. Its foundation, is a conceptual model: the "Politically Motivated Reasoning Paradigm" (PMRP). PMRP identifies the features of politically motivated reasoning that distinguish it not only from a truth-convergent Bayesian model of information processing but also from various other, non-Bayesian cognitive biases.

The validity of study designs used to test hypotheses about politically motivated reasoning depends on how readily they enable manipulation and observation of the key elements of PMRP. This essay also describes an experimental setup—the *PMRP design*—geared toward these ends. How the PMRP model and design can be used to address unresolved research questions will be the focus of the companion essay.

PMRP: A CONCEPTUAL MODEL

Motivated reasoning refers to the tendency of individuals to unconsciously conform assessment of factual information to some goal *collateral* to assessing its truth. Such goals are myriad: maintaining a positive self-conception (Dunning, 2003); rationalizing self-serving behavior (Hsee, 1996); perceiving coherence rather than complexity in evidence informing important decisions

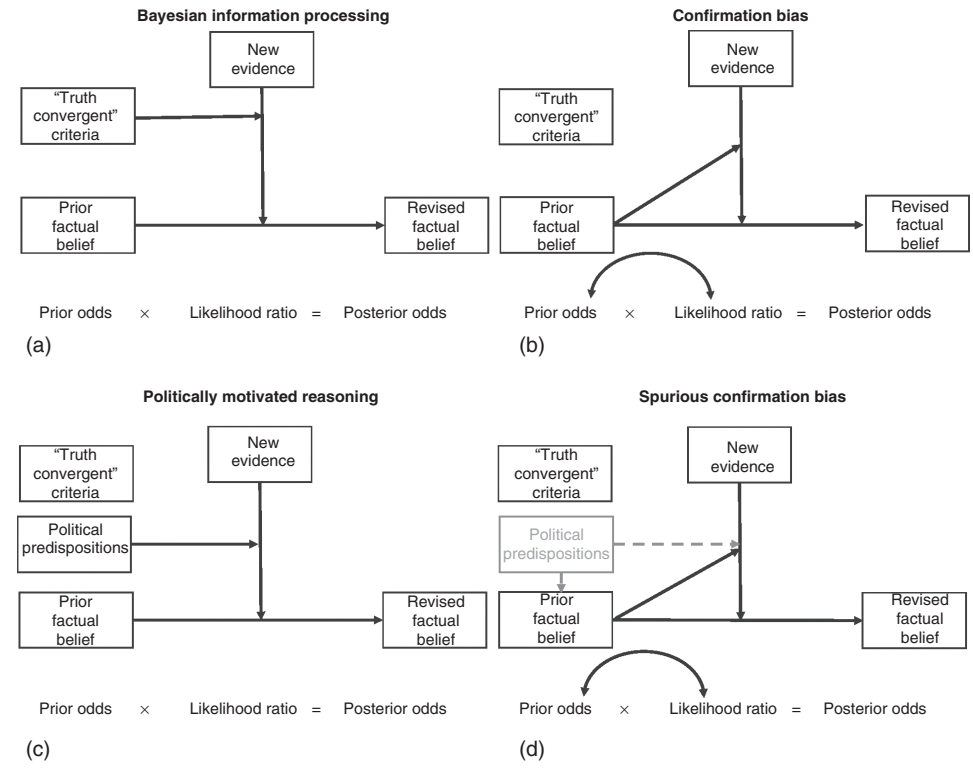


Figure 1 Models of information processing.

(Russo, Carlson, Meloy, & Yong, 2008). The truth-independent goal of “politically motivated reasoning” is *identity protection*: the formation of beliefs that maintain a person’s status in affinity group united by shared values (Cohen, 2003; Greene, 2013; Sherman & Cohen, 2006; Westfall, Van Boven, Chambers, & Judd, 2015).

It will help to begin with a clearer picture of what a “truth convergent” form of information processing looks like. Consider a barebones Bayesian model (Figure 1a). It consists of a “prior” or existing estimate of the probability of some hypothesis; a piece of new information or evidence; and a revised estimate that reflects the probative weight of that information. The weight takes the form of a *likelihood ratio*, which reflects how much more consistent the information is with the hypothesis than with some alternative. Bayes’ theorem instructs the person whose goal is to form the best revised estimate of the probability of a hypothesis to multiply her prior assessment of the probability (expressed in odds) by the information’s likelihood ratio (Lempert, 1977).

We can understand the character of any non-truth-convergent information-processing mechanism by assessing how it relates to this model. *Confirmation*

bias (Rabin & Schrag, 1999), for example, involves selectively crediting information conditional on its consistency with one's existing beliefs (Figure 1b). For example, someone who is convinced that human-caused climate change is not happening might infer from the contrary view of the National Academy of Sciences that NAS members have no expertise on this issue, and thus dismiss an Academy "expert consensus" report as unentitled to weight. In Bayesian terms, he is *deriving* the likelihood ratio for the report from his priors (Stanovich, 2011).

There is nothing in *Bayes' theorem* that forbids this. Bayes' theorem does not say *how* to figure out the likelihood ratio, only what to *do* with it: treat it as the factor by which one multiplies one's prior odds. Indeed, deriving the likelihood ratio for new information from one's priors—consciously or otherwise—might be a "sensible" strategy for maximizing one's welfare where the expected benefit of correcting a mistaken belief exceeds the cost of evaluating new information on some basis independent of one's preexisting views (Gerber & Green, 2013).

However, in order for Bayesianism to be "truth convergent," one must derive the likelihood ratio assigned new information in a truth-convergent fashion—a feature we can add to our (still) spare Bayesian model (Figure 1a). Confirmation bias is indisputably *not* truth convergent: someone who engages in it will necessarily fail to correct a mistaken perception of facts even when furnished valid contrary information (Rabin & Schrag, 1999).

Figure 1c shows how politically motivated reasoning relates to the Bayesian model. In effect, someone engaged in motivated reasoning derives the likelihood ratio for new information *not* from truth-convergent criteria independent of her priors but from the impact crediting it will have on aligning her beliefs with those of others in an identity-defining group.

Consider a study of how politically motivated reasoning can affect perceptions of *scientific consensus* (Kahan, Jenkins-Smith, & Braman, 2011). In it, members of the general public were shown highly credentialed scientists. Subjects were asked to indicate how strongly they disagreed or agreed that each scientist was an *expert* on a particular societal risk—global warming, nuclear wastes, or gun control. The *positions* of the scientists were manipulated: half the subjects believed the featured scientist held the "high risk" position, and half the "low risk" one, on the indicated issue. The subjects' assessment of the expertise of each scientist was highly correlated with whether the position attributed to the scientist matched the predominant one among individuals with the subjects' cultural outlooks.

The beliefs of any given expert scientist on a matter within that scientist's domain is one piece of evidence—about both what the relevant facts are and what experts believe those facts to be. By adjusting their assessment of whether a particular scientist was in fact an "expert" based on the position he




Figure 2 Politically motivated cognition of evidence of science consensus. Colored bars reflect 0.95 CIs. Source: Adapted from Kahan *et al.* (2011).

was represented as taking, study subjects effectively assigned this evidence a likelihood ratio equal to or greater than one depending on whether it supported or contradicted a conclusion congenial to their identities: namely, expert scientists agree with the position that predominates in *my* cultural group (Figure 2).

If individuals reason this way outside the lab, groups who are polarized on the contribution of human activity to climate change, the safety of deep geologic isolation of nuclear wastes, and the impact of concealed-carry laws on crime should hold opposing perceptions of *scientific consensus* on these issues as well. And they do (Kahan *et al.*, 2011).

This form of motivated reasoning is *not* confirmation bias, although it can easily be confused with it. Someone who engages in politically motivated reasoning will predictably form beliefs consistent with the position that fits her predispositions. Because she will also selectively credit new information based on its congeniality to that same position, it will *look* like she is deriving the likelihood ratio from her priors. However, the correlation is spurious: a “third variable”—her motivation to form beliefs congenial to her identity—is the “cause” of *both* her priors *and* her likelihood ratio assessment (Figure 1d).

This difference matters. Imagine we construct an experiment that changes subjects’ perception of how evidence relates to their political commitments. For example, we might furnish individuals’ information that manipulates the perception of how crediting evidence on climate change coheres with their groups’ identity-expressive attitudes toward free markets. We can then measure the significance subjects afford new evidence on climate change. If we assume a person will weight such information consistent with her priors due to “confirmation bias,” we should not expect the manipulation to matter.

But if we believe that someone is engaged in “politically motivated reasoning,” we should expect the manipulation to *counteract* the person’s biased determination of the likelihood ratio, because we will have altered what we hypothesize to be the *cause* of the person’s biased information processing. Indeed, in such a study, individuals whose promarket sensibilities predisposed them toward climate-change skepticism rated the strength of evidence for global warming much more highly after being exposed to information on *geoengineering*—a technological “fix” that would obviate the need for commerce-inhibiting CO₂-emission regulations. Likewise, individuals whose antimarket sensibilities predisposed them to credit evidence on climate change treated such evidence as *less convincing* after learning of geoengineering research than when they were briefed instead on the need for stronger CO₂-emission standards (Kahan, Hank, Tarantola, Silva, & Braman, 2015).

Or imagine we expose subjects to information on an unfamiliar issue—say, the risks of nanotechnology. Because by design, they lack any previous position, it’s unclear what an expectation of “confirmation bias” would predict. However, it certainly does not furnish us with reason to expect a strong relationship between post-information-exposure beliefs and subjects’ political commitments. Alternatively, if we think such individuals will react to the information by forming affective reactions that express their groups’ predispositions toward technological risks generally, we might predict information-exposed subjects will be politically polarized relative to information-unexposed ones. Experiments support this prediction (Druckman & Bolsen, 2011; Kahan, Braman, Slovic, Gastil, & Cohen, 2009).

Many researchers conflate “politically motivated reasoning” and “confirmation bias.” They should not. There is a genuine difference and study designs that fail to unconfound the two impede explanation, prediction, and prescription.

MEASURING PMRP

Experimental Data

The most compelling way to test the hypothesis that “fact polarization” originates in politically motivated reasoning is through an experiment crafted to elicit reasoning consistent with *it* and *inconsistent* with Bayesian or other forms of information processing. Designing such an experiment, however, is not straightforward.

Imagine a researcher measures subjects’ “beliefs” in “anthropogenic global warming” (AGW), and the strength of those beliefs (measured as probability assessments), and then furnishes them a study presenting evidence in

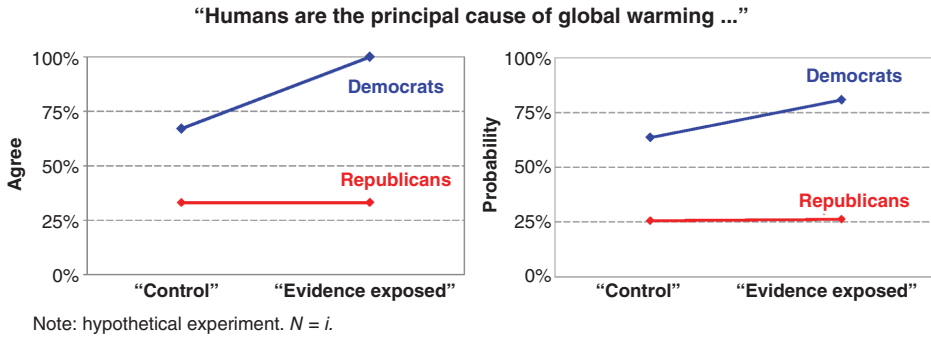


Figure 3 Hypothetical politically motivated reasoning experiment.

favor of AGW. After subjects read the study, the gap in the proportions of Democrats and Republicans who accept AGW, and in their respective estimates of the probability of it, have both widened (Figure 3).

Is this evidence of politically motivated reasoning? It’s impossible to say!

The distinctive feature of “politically motivated reasoning” is the disposition to derive the likelihood ratio for new information from one’s political predispositions rather than from truth-convergent criteria. If that happened here, we would expect to see what the researcher did: partisans becoming more “polarized” as they examine the “same” evidence. However, in fact we *could* have seen the same pattern if the subjects were assessing the information consistent with the Bayesian model (Figure 1a), too.

Assume the sample consisted of Rita, Ron, and Rose—all Republicans—and Donny, Dave, Daphne—all Democrats. Their “beliefs” about “human-caused climate change” are reflected in the “before” column of Table 1. When shown the AGW study, *they all agree about the weight it should be given*. They all perceive, let’s posit, that the study evidence has modest weight—a likelihood ratio of 3, meaning it is three times more consistent with the hypothesis that humans are responsible for climate change. In other words, *none of the subjects* adjusts the weight afforded the evidence to fit his or her predispositions.

All subjects agree that the evidence is three times more consistent with the hypothesis that “human activity is the principal cause of climate change” than with the rival hypothesis that it is not. However, because the information is not “new” for Ron and Rose—that is, has been encountered by them and assimilated to their priors before they entered the experiment—they do not revise their assessment of the probability of the rival hypotheses.

Nevertheless, this experiment could *still* display the observed polarization. *First*, the subjects started with *heterogeneous priors* (Gerber & Green, 2013). Daphne, for example, put the probability that humans are causing climate change at 0.5:1 in favor before she saw the AGW study. Rita’s prior odds

Table 1
Hypothetical-Experiment Subjects' Beliefs before and after
Information Exposure

	Before		Weight of study evidence		After	
	Prior odds on AGW	%	LR_W	LR_E	Post odds on AGW	%
Donny	2 : 1	67	3	3	6 : 1	86
Dave	10 : 1	91	3	3	30 : 1	97
Daphne	0.5 : 1	33	3	3	1.5 : 1	60
Rita	0.01 : 1	1	3	3	0.03 : 1	3
Ron	0.1 : 1	9	3	1	0.1 : 1	9
Rose	2 : 1	67	3	1	2 : 1	67

Study info refers to the evidence on climate change supplied to subjects. LR_W refers to the “likelihood ratio” or weight subjects assign that information *the first time they encounter it in the world*. LR_E refers to the weight they assign the information in the experiment.

were 0.01 : 1 in favor. When they both afforded the study a likelihood ratio of 3 in the experiment (as reflected in the LR_E column), Daphne *flipped* from the view that humans “probably” are not responsible for climate change to the view that they probably *are* (1.5 : 1 or 3 : 2 in favor). However, because Rita was more strongly convinced that humans were not causing climate change, she *persisted* in her belief *against* AGW even after appropriately adjusting her confidence level downward.

If the “outcome variable” of the study, then, is “percentage of Republicans and Democrats who think humans are causing global warming,” we will see polarization *even with* Bayesian information processing—that is, even *without* the selective crediting of information that is the signature of politically motivated reasoning (Bullock, 2009).

Second, the subjects started with differing amounts of knowledge. As it happens, Ron and Rose *already knew* about the study the researcher supplied. They assigned the study a likelihood ratio of 3 the first time they encountered it in the world (LR_W). However, their priors—10 : 1 against and 2 : 1 in favor of human-caused climate change, respectively—already reflected their unbiased assessment of it. Accordingly, when they were shown the study in the experiment (LR_E), they assigned it a likelihood ratio of “1”—*not* because they were conforming the likelihood ratio to their predispositions but because for them the study was not *new* information.

The partisan differential in the “mean” probability assigned to AGW grew (Figure 3). However, that was not a consequence of the varying weight subjects of opposing groups assigned the evidence. Rather, the source was a *pre-treatment difference in exposure to information* (Druckman, Fein, & Leeper, 2012):

because fewer Democratic subjects had previously encountered information equivalent to that supplied by the researcher, a greater proportion of them revised upward their assessment of the probability of AGW.

In sum, to draw confident inferences that politically motivated reasoning is generating polarization, we need a better study design, one that avoids the confounds of *heterogeneous priors* and *pretreatment information exposure*. The classic experiment by Lord, Ross, and Lepper (1979) comes very close.

In it, LRL (Lord, Ross, and Lepper) furnished subjects of opposing beliefs short summaries of studies on the deterrent efficacy of capital punishment. The study methods were held constant and only the represented outcomes manipulated. Effectively, then, those on both sides were furnished evidence of *equal weight*. Nevertheless, LRL reported that subjects' ratings of the strength of the studies matched their prior views. Indeed, after examining the studies, subjects of opposing positions indicated that their confidence in their positions had increased, a dynamic that LRL labeled "biased assimilation and polarization."

Even well-designed studies leave room for uncertainty. One objection to LRL's is that it lacked any objective measure of "polarization," relying only on subjects' post-information representations that their beliefs had intensified. Even more tellingly, the *within-subjects* design created reason to question the genuineness of the reported biased information processing: subjects' views on a contentious issue having been solicited immediately before being exposed to counter-attitudinal information, they might have felt constrained to deny being influenced to rebut any inference that their preexisting opinions were uninformed (Druckman, 2012; Gerber & Green, 1999).

Some more recent studies reflect a design responsive to these objections (Bolsen, Druckman, & Cook, 2014; Scurich & Shniderman, 2014; Uhlmann, Pizarro, Tannenbaum, & Ditto, 2009). Like LRL, these experiments do not rely on distinct sets of "pro-" and "con-" evidence but instead manipulate the identity-protective stake subjects have in crediting *one and the same piece of evidence*. Unlike LRL, however, these studies use *between-subject* designs: what's compared is not individual subjects' reported beliefs before and after being exposed to information but rather the weight diverse *groups* give the evidence in each experimental condition.

Such studies strongly corroborate LRL's findings. In one, subjects examined a film of protestors alleged to have physically harassed passersby (Kahan, Hoffman, Braman, Evans, & Rachlinski, 2012). The identity of the demonstrators was manipulated: in one condition, they were described as "antiabortion protestors"; in another, as "gay-rights advocates."

Subjects of opposing "cultural worldviews" who were assigned to the same condition—and who thus believed they were watching the same type of protest—reported forming *opposing* perceptions of key facts (e.g., whether

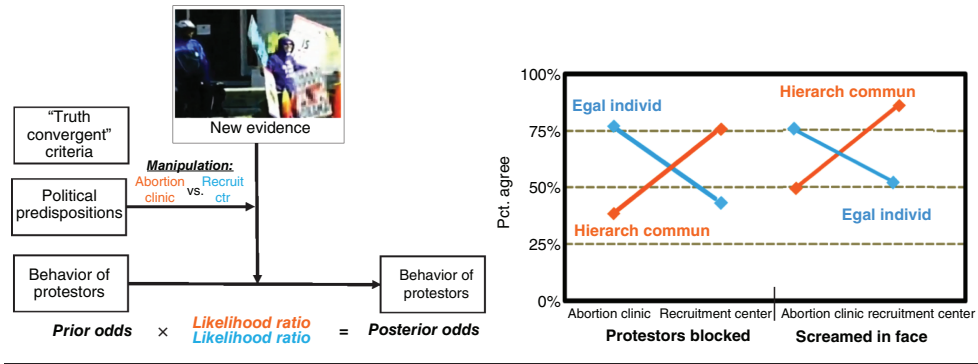


Figure 4 Politically motivated cognition of behavior of political protestors. In the PMRP Design, the identity-protective stake subjects have in crediting or discrediting *one and the same piece of evidence* is manipulated. Here, perceptions of the behavior of protestors varied depending on whether their perceived affirmed or threatened subjects’ identities. *Source:* Adapted from Kahan *et al.* (2012).

the protestors “blocked” and “screamed” at pedestrians). At the same time, subjects assigned to different conditions—who believed they were watching different types of protests—formed perceptions contrary to subjects who *shared* their worldviews (Figure 4).

Another important feature of this design is their use of subjects’ *group identities* to predict differences in assessments of evidence. By virtue of this feature, these studies supply more secure grounds than do ones that characterize subjects based on *their prior positions* for attributing the results to *motivated reasoning* rather than confirmation bias.

I will call this experimental setup—a between-subject one that assesses the *weight* assigned a single piece of evidence conditional on experimental manipulation of its perceived identity congruence—the *PMRP design*. The PMRP design is *not* the only one that validly measures politically motivated reasoning. Indeed, the consistency of findings of studies that reflect the PMRP design and ones based on alternatives (Druckman & Bolsen, 2011; Nyhan, Reifler, Richey, & Freed, 2014) furnishes more reason to credit both. Nevertheless, the *test* the PMRP design is constructed to pass—demonstration that individuals are *adjusting* the weight assigned evidence conditional on its identity congruence—supplies the proper standard for assessing whether any particular study design supports an inference of politically motivated reasoning.

Whether a study that reflects the PMRP design should be deemed to require evidence of *changes* in the subjects’ “beliefs” (their “posteriors,” in Bayesian terms) should depend on the strength of the inferences the study otherwise supports. In mock-juror experiments (Kahan *et al.*, 2012; Scurich & Shniderman, 2014) and novel-issue ones (Druckman & Bolsen, 2011; Kahan *et al.*,

2009), subjects will have had no meaningful prior exposure to evidence on the contested facts. Because subjects' "beliefs" are thus artifacts of the experiment, one would expect the manipulation to generate mirror-image states of belief polarization between the experimental conditions.

The situation is different, however, in experiments examining familiar policy disputes (climate change, gun control, etc.). On these, many subjects will start with strongly held views based on information encountered before the experiment. The failure of the manipulation to generate "changed beliefs" does not rule out an inference of politically motivated reasoning under such circumstances for exactly the same reason that persistence of beliefs among subjects exposed to counter-attitudinal evidence does not *support* such an inference: the unobserved differential in the strength of the subjects' "priors" will confound *any* inference that they evaluated the strength of the study evidence in a manner unaffected by its congeniality to their predispositions.

What matters is whether the design enables a confident assessment that individuals of opposing outlooks adjusted the weight assigned evidence in response to the experimental manipulation. If so, their reasoning displayed the essential feature of PMRP. By virtue of the strength of their priors, they might not have altered their *beliefs* on the issue in question. However, if the experiment in fact captured how individuals assess comparable information outside the lab, then, contrary to what would occur under the Bayesian model, individuals of opposing outlooks will not converge no matter how much valid evidence they are furnished. Or will not unless something is done to change the identity-protective stake individuals have in forming those beliefs.

Operationalizing Identity

Scholars use diverse frameworks to measure the predispositions that inform politically motivated reasoning. Left-right political outlooks are most common (Lodge & Taber, 2013). "Cultural worldviews" are used in studies (Bolsen, Druckman, & Cook, 2015; Druckman & Bolsen, 2011) that investigate "cultural cognition," a construct developed to explore societal-risk conflicts (Kahan, 2015).

The question whether politically motivated reasoning is "really" driven by "ideology" or "culture" or some other source of affinity is ill-posed. One might take the view that myriad commitments—including not only political and cultural outlooks but religiosity, ethnicity, gender, region of residence, and so on—figure in politically motivated reasoning on "certain occasions" or to "some extent." Much better, however, would be to recognize that *none* of these characteristics is the "true" source of the predispositions that inform politically motivated reasoning. All are simply imperfect proxies for an unobserved shared disposition that orients information processing.

Studies that use alternative predisposition constructs, then, are not testing alternative *theories* of “what” motivates politically motivated reasoning. They are simply employing alternative *measures* of whatever it is that does.

The only reason for preferring one measure over another is explanatory, predictive, and prescriptive utility. The best test of whether a researcher is using the “right” one is what she is able to *do* with it.

Samples

Early politically motivated reasoning experiments were conducted on students, but today studies use a diverse variety of subjects. Does sample composition matter?

The answer, here as elsewhere, depends on whether the sample supports the inference that a researcher is drawing (Druckman & Kam, 2011). Samples constructed to represent the general population normally will. However, so will *nonrepresentative* convenience samples so long as, *first*, they contain a sufficient number of the varying types of subjects with the identity-protective stake hypothesized to generate factual polarization in the real world; and *second*, the subjects are *typical* of the real-world people the study is model.

Consider a sampling strategy targeting persons who visit an Internet site because they are interested in moral psychology. Studies of such individuals can certainly enlarge general understanding (Iyer, Koleva, Graham, Ditto, & Haidt, 2012). However, because people who have this particular interest are highly unusual, such a sampling method will not support confident inferences about how individual differences in reasoning relate to opposing political outlooks in the general population.

Student samples could be valid. However, they obviously will not be if they do not contain individuals who possess the identities of the general-population groups whose members are polarized. For example, because being white, being male, and being conservative are all important indicators of the cultural identity with the greatest stake in resisting claims of environmental risks (McCright & Dunlap, 2013), a study sample that comprises “forty-one New York University undergraduates (30 women, 11 men)” (Feygina, Jost, & Goldsmith, 2010) will not plausibly support valid inferences about “messages” likely to offset politically motivated reasoning among “climate change skeptics.”

Under the proposed criteria, studies using M Turk workers (who on average participate in 300 “studies” over the course of their “careers,” with some having participated in thousands (Rand *et al.*, 2014)) are not valid. Such samples tend to be ideologically skewed toward the left (Richey & Taylor, 2012). While such a deficiency could be remedied by oversampling “conservatives,” the existence of the ideological imbalance in the M Turk

workforce creates reasons to be suspicious about the *typicality* of those “workers” who report holding a “conservative” ideology: if the tasks M Turk workers typically perform [such as tagging online pornography (Dobson, 2013)] deter ordinary conservatives from participating in the M Turk workforce, for example, the self-identified conservatives who choose to join it nonetheless are likely to have sensibilities unlike those in the general population. Indeed, the demonstrated propensity of non-US M Turk workers to use US-VPNs to disguise their nationality (Shapiro, Chandler, & Mueller, 2013) and engage in like forms of misrepresentation to qualify for studies (Chandler & Shapiro, 2016) furnishes reason to doubt M Turk workers are “typical” of ordinary members of the US population generally. Not surprisingly, at least one study has shown M Turk workers do *not* respond comparably to general population or student samples in politically motivated reasoning studies (Krupnikov & Levine, 2014).

NOW WHAT?

Whether politically motivated reasoning best explains fact polarization remains strongly debated. Moreover, even among scholars who believe the weight of existing evidence supports such a conclusion, there remain myriad issues of disagreement. The contribution the PMRP model and design can make to examination of such questions is the subject of the second in this two-essay set.

REFERENCES

- Bolsen, T., Druckman, J. N., & Cook, F. L. (2014). The influence of partisan motivated reasoning on public opinion. *Political Behavior*, *36*, 235–262.
- Bolsen, T., Druckman, J. N., & Cook, F. L. (2015). Citizens’, scientists’, and policy advisors’ beliefs about global warming. *The ANNALS of the American Academy of Political and Social Science*, *658*(1), 271–295.
- Bullock, J. G. (2009). Partisan Bias and the Bayesian Ideal in the Study of Public Opinion. *Journal of Politics*, *71*, 1109–1124.
- Chandler, J., & Shapiro, D. (2016). Conducting clinical research using crowdsourced convenience samples. *Annual Review of Clinical Psychology*, *12*, 53–81.
- Cohen, G. L. (2003). Party over policy: The dominating impact of group influence on political beliefs. *Journal of Personality and Social Psychology*, *85*, 808–822.
- Dobson, J. (2013). Mechanical Turk: Amazon’s New Underclass. Huffington Post, http://www.huffingtonpost.com/julian-dobson/mechanical-turk-amazons-underclass_b_2687431.html
- Druckman, J. N. (2012). The politics of motivation. *Critical Review*, *24*, 199–216.

- Druckman, J. N., & Bolsen, T. (2011). Framing, motivated reasoning, and opinions about emergent technologies. *Journal of Communication*, 61(4), 659–688. doi:10.1111/j.1460-2466.2011.01562.x.
- Druckman, J. N., Fein, J., & Leeper, T. J. (2012). A source of bias in public opinion stability. *American Political Science Review*, 106, 430–454.
- Druckman, J. N., & Kam, C. D. (2011). Students as experimental participants. In J. N. Druckman, D. P. Green, J. H. Kuklinski & A. Lupia (Eds.), *Cambridge handbook of experimental political science* (pp. 41). New York, NY: Cambridge University Press.
- Dunning, D. (2003). The relation of self to social perception. In M. R. Leary & J. P. Tangney (Eds.), *Handbook of self and identity* (pp. 421–441). New York, NY: Guilford Press.
- Feygina, I., Jost, J. T., & Goldsmith, R. E. (2010). System justification, the denial of global warming, and the possibility of “System-Sanctioned Change”. *Personality & Social Psychology Bulletin*, 36, 326–338.
- Frankovic, K. *Most republicans do not think humans are causing climate change*. YouGov. (2015), <https://today.yougov.com/news/2015/11/25/most-republicans-do-not-think-humans-causing-clima/>.
- Greene, J. D. (2013). *Moral tribes: Emotion, reason, and the gap between us and them*. New York: Penguin Press.
- Gerber, A., & Green, D. (1999). Misperceptions about perceptual bias. *Annual Review of Political Science*, 2, 189–210.
- Hsee, C. K. (1996). Elastic justification: How unjustifiable factors influence judgments. *Organizational Behavior and Human Decision Processes*, 66, 122–129.
- Iyer, R., Koleva, S., Graham, J., Ditto, P., & Haidt, J. (2012). Understanding libertarian morality: The psychological dispositions of self-identified libertarians. *PLoS ONE*, 7, e42366.
- Jost, J. T., Hennes, E. P., & Lavine, H. (2013). “Hot” political cognition: Its self-, group-, and system-serving purposes. In D. E. Carlson (Ed.), *Oxford handbook of social cognition* (pp. 851–875). New York, NY: Oxford University Press.
- Kahan, D. M. (2015). Climate-science communication and the measurement problem. *Advances in Political Psychology*, 36, 1–43.
- Kahan, D. M. (in press). The expressive rationality of inaccurate perceptions of fact. *Behavioral & Brain Sciences*.
- Kahan, D. M., Braman, D., Slovic, P., Gastil, J., & Cohen, G. (2009). Cultural cognition of the risks and benefits of nanotechnology. *Nature Nanotechnology*, 4, 87–91.
- Kahan, D. M., Hank, J.-S., Tarantola, T., Silva, C., & Braman, D. (2015). Geoengineering and climate change polarization: Testing a two-channel model of science communication. *Annals of the American Academy of Political and Social Science*, 658, 192–222.
- Kahan, D. M., Hoffman, D. A., Braman, D., Evans, D., & Rachlinski, J. J. (2012). They saw a protest: Cognitive illiberalism and the speech-conduct distinction. *Stanford Law Review*, 64, 851–906.
- Kahan, D. M., Jenkins-Smith, H., & Braman, D. (2011). Cultural cognition of scientific consensus. *Journal of Risk Research*, 14, 147–174.

- Krupnikov, Y., & Levine, A. S. (2014). Cross-sample comparisons and external validity. *Journal of Experimental Political Science*, 1, 59–80.
- Lempert, R. O. (1977). Modeling Relevance. *Michigan Law Review*, 75, 1021–1057.
- Lodge, M., & Taber, C. S. (2013). *The rationalizing voter*. Cambridge, MA: Cambridge University Press.
- Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization. *Journal of Personality and Social Psychology*, 37, 2098–2109.
- McCright, A. M., & Dunlap, R. E. (2013). Bringing ideology in: The conservative white male effect on worry about environmental problems in the USA. *Journal of Risk Research*, 16, 211–226.
- Moore, P. (2015). *Most Republicans back a flat tax*. YouGov, <https://today.yougov.com/news/2015/05/22/flat-tax/>
- Newport, F. *Majority say more concealed weapons would make U.S. Safer*. Gallup 2015, <http://www.gallup.com/poll/186263/majority-say-concealed-weapons-safer.aspx>
- Nyhan, B., Reifler, J., Richey, S., & Freed, G. L. (2014). Effective messages in vaccine promotion: A randomized trial. *Pediatrics*, 133, e835–e842.
- Rabin, M., & Schrag, J. L. (1999). First impressions matter: A model of confirmatory bias. *Quarterly Journal of Economics*, 114, 37–82.
- Rand, D. G., Peysakhovich, A., Kraft-Todd, G. T., Newman, G. E., Wurzbacher, O., Nowak, M. A., & Greene, J. D. (2014). Social heuristics shape intuitive cooperation. *Nature Communication*, 5, 3677.
- Richey, S., & Taylor, B. (2012). How representatives are Amazon mechanical turk workers? *The Monkey Cage*, <http://themonkeycage.org/2012/12/how-representative-are-amazon-mechanical-turk-workers/>
- Russo, J. E., Carlson, K. A., Meloy, M. G., & Yong, K. (2008). The goal of consistency as a cause of information distortion. *Journal of Experimental Psychology*, 137, 456.
- Scurich, N., & Shniderman, A. B. (2014). The selective allure of neuroscientific explanations. *PLoS One*, 9, e107529.
- Shapiro, D. N., Chandler, J., & Mueller, P. A. (2013). Using mechanical turk to study clinical populations. *Clinical Psychological Science*, 1, 213–220.
- Sherman, D. K., & Cohen, G. L. (2006). The psychology of self-defense: Self-affirmation theory. *Advances in Experimental Social Psychology*, 38, 183–242.
- Smith, T.W., Marsden, P.V., & Hout, M. *General social survey, 2014 file [data file and codebook]*. Chicago, IL: National Opinion Research Center. Ann Arbor, MI: Inter-university Consortium for Political and Social Research (2014), http://gss.norc.org/Documents/spss/2014_spss.zip
- Stanovich, K. E. (2011). *Rationality and the reflective mind*. New York: Oxford University Press.
- Uhlmann, E. L., Pizarro, D. A., Tannenbaum, D., & Ditto, P. H. (2009). The motivated use of moral principles. *Judgment and Decision Making*, 4, 479.
- Westfall, J., Van Boven, L., Chambers, J. R., & Judd, C. M. (2015). Perceiving Political Polarization in the United States Party Identity Strength and Attitude Extremity Exacerbate the Perceived Partisan Divide. *Perspectives on Psychological Science*, 10, 145–158.

DAN M. KAHAN SHORT BIOGRAPHY

Dan M. Kahan is the Elizabeth K. Dollard professor of law and professor of psychology at Yale Law School. His primary research interests (for the moment, anyway) are risk perception, science communication, and the application of decision science to law and policymaking. He is a member of the Cultural Cognition Project, an interdisciplinary team of scholars who use empirical methods to examine the impact of group values on perceptions of risk and related facts. In studies funded by the National Science Foundation, his research has investigated public disagreement over climate change, public reactions to emerging technologies, and conflicting public impressions of scientific consensus.

RELATED ESSAYS

Models of Revealed Preference (*Economics*), Abi Adams and Ian Crawford
 Economic Models of Voting (*Political Science*), Ian G. Anson and Timothy Hellwig

Public Opinion and International Conflict (*Political Science*), Adam J. Berinsky

Coevolution of Decision-Making and Social Environments (*Sociology*), Elizabeth Bruch *et al.*

Mental Models (*Psychology*), Ruth M. J. Byrne

Choice Architecture (*Psychology*), Adrian R. Camilleri and Rick P. Larrick

Culture and Cognition (*Sociology*), Karen A. Cerulo

Misinformation and How to Correct It (*Psychology*), John Cook *et al.*

Cognitive Processes Involved in Stereotyping (*Psychology*), Susan T. Fiske and Cydney H. Dupree

Political Advertising (*Political Science*), Erika Franklin Fowler

Emotion and Decision Making (*Psychology*), Jeff R. Huntsinger and Cara Ray

Genetic Foundations of Attitude Formation (*Political Science*), Christian Kandler *et al.*

Reconciliation and Peace-Making: Insights from Studies on Nonhuman Animals (*Anthropology*), Sonja E. Koski

From Individual Rationality to Socially Embedded Self-Regulation (*Sociology*), Siegwart Lindenber

Implicit Attitude Measures (*Psychology*), Gregory Mitchell and Philip E. Tetlock

Event Processing as an Executive Enterprise (*Psychology*), Robbie A. Ross and Dare A. Baldwin

Stereotype Threat (*Psychology*), Toni Schmader and William M. Hall