

## **Why Swing State Voting Is Not (Usually) Effective Altruism: The Bad News about the Good News about Voting**

Conventional wisdom in political science long held that the odds of casting a decisive vote in a large-scale election are extraordinarily small; thus, the expected social value of a vote is effectively zero.<sup>1</sup> An altruistic citizen would do more expected good spending election day raising funds for the Against Malaria Foundation than by pulling a lever at the polls.

But many philosophers and social scientists have recently defended what we will call the *Charitable Voting Hypothesis*: casting a vote in a key swing state is similar to donating thousands of dollars to charity. Aaron Edlin, Andrew Gelman, and Noah Kaplan write, “At \$50,000 per half hour, voting is surely one of the best and most efficient charities around for a voter in New Mexico.”<sup>2</sup> William MacAskill argues that for certain citizens “voting is like donating thousands of dollars to (developed-world) charities . . . That’s a much better use of your time than you could get...by working the hour it takes you to vote and donating your earnings.”<sup>3</sup> MacAskill in particular defends swing state voting as a form of effective altruism—the practice of allocating one’s philanthropic resources in ways that maximize social welfare gains. Zach Barnett has recently argued that voting is rational in sufficiently close high-stakes elections if one has public-spirited goals.<sup>4</sup> The Charitable Voting Hypothesis analogizes a vote to a ticket to a lottery with a huge cash prize. The odds of winning are low, but the expected value is high enough to be worth buying.<sup>5</sup>

We contend that the argument for the Charitable Voting Hypothesis is overstated. For voting to be a high expected utility activity, it’s not enough to cast a vote that has a sufficiently high chance of deciding the election in a high-stakes election. One must also cast a vote for the higher utility candidate. Defenders of the Charitable Voting Hypothesis may have succeeded in showing that the odds of casting a decisive vote are higher than previously thought; however,

this only takes them half of the way to a vindication of voting as a high expected utility activity. The second half of the way is the hard part.

In this paper, we argue that the Charitable Voting Hypothesis at most applies only to a very small number of unusually epistemically privileged citizens. If votes matter, they can have high expected social benefits but also equivalently high social costs. A publicly-spirited citizen should thus take voting seriously and vote only if they are justified in believing they are voting for the higher utility candidate. However, most citizens—*especially* those who are confident they are voting well—lack such justification. Even in swing states in close and high stakes elections, the majority of voters should not regard their votes as being equivalent to donating thousands of dollars to charity.

We begin by summarizing recent developments in the literature on the decisiveness of a vote (§1). Next we explain that defenders of the Charitable Voting Hypothesis face an unmet challenge—viz., determining which candidate is higher utility (§2). We then argue that problems of inaccurate forecasting, general cognitive biases, partisan political biases, assessments of counterfactual presidencies, and normative uncertainty make it exceedingly difficult to form justified beliefs about the utility difference between two candidates (§3). Moreover, the cost of attempting to overcome these epistemic obstacles is likely greater than the benefit (§4). We conclude that voting remains a low expected value activity for most citizens (§5).

### **1. The Expected Utility of a Vote: The Basic Model**

Suppose Bill plans to vote. Suppose Bill's deciding to vote has no effect on politics other than possibly changing the election's outcome. Suppose that Bill derives no consumption value

from voting. Suppose there are two candidates, D and R. If Bill decides to vote for D instead of R, the expected value of his vote ( $U_v$ ) is given by equation (1)<sup>6</sup>:

$$(1) U_v = p_d \Delta_c - C_o$$

where

- $p_d$  represents the probability that Bill's will cast the decisive vote,
- $\Delta_c$  represents the difference in the expected value of the two candidates, in terms of whatever consequence they will produce if elected. That is,  $\Delta_c = U_D - U_R$ .
- $C_o$  represents Bill's opportunity cost—the next best thing Bill could have done with his time.

The expected utility of Bill's vote is the difference in the expected value between his favored candidate over the other candidate, discounted by the probability that Bill's vote will be decisive, minus his opportunity cost.<sup>7</sup> Note that  $U_v$  will be negative if Bill has a high opportunity cost or if Bill votes for the worse candidate.

Voting is somewhat like buying a lottery ticket. Lottery tickets often have high winning payoffs but low chances of winning. If  $p_d \Delta_c < C_o$ , then it is irrational for Bill to vote, and he should instead perform his next best activity.

However, voting differs from lotteries in that a losing lottery ticket pays \$0, while a decisive bad vote has a social cost. This, as we will see below, often makes it harder to justify voting than to justify buying a lottery ticket.

There is significant debate among economists and political scientists over the precise way to calculate  $p_d$ , the probability that a vote will be decisive. Geoffrey Brennan and Loren Lomasky

defended a model of  $p_d$  based, roughly, on the probability that a weighted coin will land heads exactly 50% of the time.<sup>8</sup> Their model implied that the probability of being decisive is so vanishingly small that the expected utility of a single vote is negative, once we account for opportunity cost.

However, many theorists now contend that the Brennan-Lomasky model, and similar “binomial models” of voting in general, are inadequate, as such models predict elections will not be as close as we in fact find them.<sup>9</sup> MacAskill instead relies upon the Edlin, Gelman, and Kaplan (hereafter EGK) model, which accounts for state-level polling data and suggests that the chances that a vote will be decisive are much higher than what is suggested by the Brennan-Lomasky model. Without getting into the details about how they estimate  $p_d$ , Edlin, Gelman, and Kaplan (hereafter EGK) argue that an American voter in a presidential election could have something like a 1 in 10 million chance in being decisive in certain close states. EGK think that most voters have a far lower chance, some on the order of 1 in a billion, while Californians have vanishingly small chances.<sup>10</sup> Both MacAskill and EGK argue that even that voting is a good idea even for many low-chance voters. If, say,  $\Delta_c = \$300$  billion, then having a 1 in a 100 million chance of being decisive means that  $p_d \Delta_c = \$3000$ . They say that since it takes only about 15 minutes to vote,  $C_o$  is low. Thus, even low-chance voters are in effect donating a few hundred or thousand dollars to charity. In a more recent paper, Barnett defends a different model for calculating the expected value of a vote which gives somewhat similar results. The fine details will not concern us here, as Barnett’s work suggests votes matter in roughly the same occasions (i.e., close-enough, high-stake elections) as the EGK model.

We assume here that either EGK or Barnett is correct about the probability a vote will be decisive or will have significant expected utility. And If EKG or Barnett are right, this seems like good news. 20 years ago, almost everyone thought that individuals' votes almost never matter (in the sense of having high expected utility), but the good news is that a far higher percentage of votes matter, even if perhaps most still do not.

However, there is bad news about this good news. The EGK or Barnett models are double-edged swords. As Jason Brennan says, showing that votes have a high chance of decisiveness simultaneously shows that some votes are good but opposing votes are bad.<sup>11</sup> The new models imply that voting well has high expected utility, but voting badly has equivalent expected disutility.

As Lomasky and Geoff Brennan note, a vote's value "must be discounted not only by the probability of being decisive but, additionally, by the probability that one has overestimated or even reversed the respective merits" of the candidates.<sup>12</sup> MacAskill also acknowledges this point: "If you're uncertain about which party is really better, you might reasonably think it's an overestimate: your expected value of voting will be lower due to a greater chance of voting for the worse party; and if you're completely unsure which party is better the expected value of voting drops to zero."<sup>13</sup> However, we argue both positive and normative uncertainty are far more significant problems than MacAskill admits. Indeed, they are problematic enough to render voting a low expected value activity for most voters even in close-enough, high-stakes elections. If the new work on probability is correct, most citizens should be wary of their own votes.

To illustrate our point, consider an analogy.<sup>14</sup> Suppose there is a purplish-red button on the other side of the state. Long thought to be purely decorative, new research suggests that the

button in fact possesses magical powers. But there's a twist: if the button is the color burgundy, pressing it will result in a \$3000 donation to an orphanage. If the button is maroon, then pressing it will destroy \$3000 of the orphanage's funds. Is driving over and pressing the button a high expected value activity?

The answer clearly depends on how capable one is of differentiating maroon and burgundy. If you know which color the button is, then push the button. But suppose there is very strong evidence that most people do not reliably distinguish burgundy from maroon. Suppose in part this is because the majority of people are ignorant of what the words "burgundy" and "maroon" mean.

Suppose also that we have strong empirical evidence that those who think they know what the words mean are nevertheless unreliable in distinguishing between them in practice. Moreover, there is also very strong evidence which shows that people who believe they are the exception to this rule—who believe they can reliably distinguish these colors—are unexceptional. That is, the evidence shows that most people would *believe* they are competent to choose even if they are not. In fact, suppose the evidence shows that people who are the most confident they can distinguish the colors are also among the least reliable. If these facts obtained, there would be powerful reasons to doubt that pressing the button is effective altruism.

We contend that this situation is analogous to the situation voters face. Actual voters suffer from all the problems we attributed to the hypothetical button-pushers, and more. Thus, even swing state voting will rarely count as effective altruism for most voters.

## **2. The Assessment of Candidates**

As noted, for your vote to be high value, your vote must have a sufficiently high chance of deciding the election *and* must be cast for the higher value candidate. We claim that forming justified beliefs about the value difference between candidates is far more difficult than defenders of the Charitable Voting Hypothesis suggest.

At first glance, MacAskill and EGK seem to describe the difference between candidates as a purely subjective assignment of expected value. For instance, EGK's argument assumes that "you *believe* that candidate A will benefit the population on average by \$1000 per person."<sup>15</sup> MacAskill spells out his reasoning in more depth:

If your preferred party is the Republicans, then you might expect to benefit because you'll pay fewer taxes. If your preferred party is the Democrats, you might expect to benefit because you'll receive more government-funded public services. Suppose for the sake of argument you conclude that your preferred party getting into power is worth \$1,000 to you. Although this \$1,000-per-citizen figure is hypothetical, it seems plausible to me.<sup>16</sup>

Interpreted literally, these passages seem to ask voters or readers to simply assign a value to the expected difference between the candidates in an election. However, the more charitable interpretation is that they intend such estimates to be more objective than that, on par with how they estimate the welfare effects of medicine, charity, or laws. Presumably, effective altruists who advocate voting intend to be consistent in their treatment of both private philanthropy and voting. In the former case, MacAskill appeals to *justified* beliefs about effectiveness. After all, the EA movement is in part a critique of philanthropic projects which people *believed* were doing good, but which are in fact ineffective or downright harmful—from the inefficient PlayPump water supply system to the Scared Straight program which results in a \$200 worth of *harm* per dollar spent.<sup>17</sup>

MacAskill denies that distributing PlayPumps or instituting Scared Straight programs are good so long as one *believes* that these programs work. Rather, effective altruists claim one should use extensive social scientific evidence to rigorously assess whether these programs in fact work. That's the very point of effective altruism. Accordingly, then, when discussing the value of a vote, the value difference between two candidates should be assessed using effective altruist's normal standards of evidence and value. If, as MacAskill claims, voting is like donating thousands to charity, then presumably he intends voters to use the same evidentiary standards and care with voting as the EA movement demands of charitable giving. Similarly, EGK presumably intend to measure the value of votes the way social scientists generally measure the welfare effects of various policies; they would not want to say that a policy counts as beneficial so long as people believe it is beneficial. In short, an effective altruist voter needs *justified* (or perhaps even *correct*) beliefs about the value difference between candidates.

In what follows, we explain why nearly everyone has grounds for being self-skeptical about their ability to judge the difference in the expected value of the two candidates. For one, we lack the information needed to actually make expected value calculations, so we must rely on, in MacAskill's terms, "guesstimation."<sup>18</sup> But political ignorance, general cognitive biases, and partisan political biases sabotage our capacity for intuiting the welfare effects of different policy bundles or to engage in accurate guesstimation. Moreover, unresolved philosophical questions about *whose* welfare counts raise serious problems of normative uncertainty.

Second, the opportunity cost of attempting to increase one's justified confidence in one's beliefs about the difference in the expected value of the two candidates is much higher than the one hour suggested by EGK and MacAskill. Such an attempt would require citizens to undertake

labor-intensive research and debiasing techniques—techniques whose effectiveness remains questionable. Thus, the cost of attempting to improve one’s vote is likely to exceed the benefit. We discuss this issue further below.

### **3. Justified Beliefs about the Value Difference Between Candidates**

Let’s start with information deficits. As a general point, it is simply hard to predict what the outcome of a given party being in power will be. The research overwhelmingly shows that even experts have a very poor track record when it comes to making specific predictions about governments and economies; experts are good at easy problems but terrible at hard problems.<sup>19</sup>

By way of initial support for the case for epistemic humility about political forecasting, consider how difficult it is to consistently make money betting on political prediction markets. You may, for instance, place bets on a variety of political matters such as Supreme Court nominations, Supreme Court decisions, GDP growth under a given presidential administration, climate change, and more. That most of us are unable to make money from these markets (and indeed, that such markets exist at all) is evidence of the unpredictability of major political events.<sup>20</sup> Indeed, if you are not able to get rich from such markets, then you have evidence that you are not a reliable predictor when it comes to politics. (One needn’t be selfish to bet; a reliable predictor who is also an effective altruist would donate winnings to effective charities.)

As difficult as it may be to predict the long-term, global effects of a *single* policy or political decision, it is far more difficult to predict the long-term, global effects of the *bundle* of policies resulting from a given presidential administration. Predictions are hard in part because, as Lomasky and Brennan note, “political uncertainties are not all of one species but several, and they compound each other.”<sup>21</sup> For instance, some of the policies likely to be instituted by

Republicans are going to be better than those likely to be instituted by the Democrats—and vice versa. But how much better are the good Republican policies than the good Democratic policies?

Or suppose that R and D differ on 10 key policies and D is better on 6 of the 10. How confident can we be that D's 6 good policies will be their priority as opposed to the 4 bad ones? Does the candidate intend to keep their campaign promises? What range of policies might D in fact implement if they change their mind or were lying? Even if D wants to make their good policies a priority, might they be politically constrained in how successfully they can pursue them? How might D react to possible but unlikely natural or political events?

Moreover, as Gerald Gaus notes, complexity undermines our ability to make accurate all-things-considered predictions of policy interventions:

Prediction becomes surprisingly difficult as soon as feedbacks enter in. But still, it might be thought, if what we really care about is only manipulating W [e.g., the employment effects of minimum wage laws], so long as there are not too many layers of feedback and we have identified many of the relevant variables, we can have a reasonable chance at successful policy interventions aiming to manipulate W's value [...] But unless we are very limited in our concerns, we would like to predict and perhaps control other aspects of the social order as well [...] The problem is now becoming clearer: as we seek to influence a number of variables—W, X, Y and Z, even if we do secure some initial leverage on W, we find that it affects, and is affected by, a host of other variables that we are also interested in manipulating. Thus...one intervention produces effects that may (or may not) give us some of what we aimed for, but which produces unexpected changes that require a policy intervention, which in turn produces effects that further require amelioration.<sup>22</sup>

Thus, it is exceedingly difficult to determine which candidate's favored policy bundle will fare better, all things considered.

Relatedly, policies that, in themselves, may have beneficial results can be in tension with one another. To take an important example, progressive redistribution and immigration liberalization are both, plausibly, welfare-enhancing policies. However, there is evidence to suggest that increased immigration reduces political support for redistribution.<sup>23</sup> There is good

evidence that the welfare effects of immigration are stronger than redistribution, though of course such findings are disputed.<sup>24</sup> Thus, the cost of implementing one welfare-enhancing policy may be the loss of a different one.

Even if one somehow knows which policies are best overall, strategic considerations complicate predictions even further. Suppose that increased immigration is a welfare enhancing policy. It may nevertheless be the case that liberalizing immigration in the short term leads to a backlash against immigration in the long term.<sup>25</sup> Perhaps citizens' attitudes toward immigration will shift toward openness as a partisan response to immigration restrictions imposed by the other side.<sup>26</sup> One may know that liberalizing immigration is welfare enhancing without knowing the best means of achieving a more liberal immigration policy.

#### *The Problem of Counterfactual Presidencies*

Consider also that defenders of the Charitable Voting Hypothesis offer, at most, *post-hoc* assessments of the impact of *particular* policies favored by presidential administrations. But to know the value difference between candidates, we need *forecasts* of the *full* effects not only of the presidential administration in power, but also an estimate of the effects of the counterfactual presidential administration that never came to power.

By way of illustration, Barnett points to the Iraq War as a high-cost result of the Bush administration, a result that was impactful enough to justify voting against Bush on act consequentialist grounds.<sup>27</sup> Yet it's worth emphasizing how difficult it would have been to predict the Iraq War in the year 2000, particularly given the role of the unexpected events of September 11<sup>th</sup>, 2001 in laying the groundwork for the war. Indeed, before 9/11, the Bush administration was concerned with Russia and China and there was little reason to expect it to invade Iraq.<sup>28</sup> It's

not as though voters could know in 2000 that Bush would launch two disastrous wars in the Middle East. To make the case for choosing Gore over Bush, we must avail ourselves of knowledge we lacked at the time *and* knowledge we lack now (how Gore would have performed).

Moreover, there is reason to think that the Bush administration brought about a number of high-benefit, but low-visibility interventions that might outweigh the harm done by its other policies. Robert Wiblin explains:

The policies which are most impactful are not always the most salient. George W. Bush's famous choice to pursue the Iraq War resulted in the removal of Saddam Hussein, though at the cost of hundreds of thousands of civilian lives and trillions of dollars in spending. But President Bush also dramatically raised US spending on antiviral drugs for impoverished victims of HIV in Africa. This 'PEPFAR' program probably would not have been pursued in his absence, and likely prevented several million deaths.<sup>29</sup>

To be clear, we take no stand on the net welfare effects of the Bush presidency. Our point is rather that these are extremely difficult to calculate after the fact, let alone before the election.

Moreover, to know the value *difference* between a Bush presidency and, say, a Gore presidency, we'd need an all-things-considered calculation of the net welfare effects of a counterfactual Gore administration. We generally do not have high-quality post-hoc assessments of the overall welfare effects of different presidents, let alone counterfactual assessments of what the losers would have done, let alone *ex ante* assessments demonstrating which candidate is really superior to another. After all, as far as we know—after doing an extensive search to write this very paper—no high-quality journal has published an estimate of the net welfare effects of an entire presidency or congressional tenure, and certainly none has published an estimate of the counterfactual welfare effects of the what the losing candidate would have done. Indeed, the one paper assessing US *state* governance on par with the rigor of GiveWell.org fails to find any

significant difference between Republican or Democratic state governance over a large variety of outcomes.<sup>30</sup>

In reply, one could argue that swing-state voting is high value even if you are merely 60 or 70 percent confident in your judgment of the higher value candidate.<sup>31</sup> But, to be consistent, we must insist that it is not enough to simply be 60 percent confident in your political judgment—your confidence must be *justified*. By analogy, defusing a nuclear bomb is a high value activity. Still, someone who is 60 percent confident they can defuse rather than detonate the bomb by cutting the red wire ought not to cut it unless they are *justified* in their confidence. The same goes for high stakes elections. And there may be voters for whom this degree of confidence is justified. But it is not enough to *assert* this confidence. Rather, one must explain how this judgment is arrived at. The burden of justification, then, rests on that would-be voter, and the evidence shows most voters lack it.

### *Normative Uncertainty*

Another obstacle to casting the right vote looms. Suppose you have a crystal ball that makes perfect predictions about the outcomes brought about by different would-be presidential administrations. You *still* shouldn't be confident in your calculations of the value difference. Normative uncertainty poses a further problem: *whose* welfare counts?

By way of example, consider the problem of abortion. Some philosophers argue that the future welfare of a fetus is as morally significant as the future welfare of an adult human being.<sup>32</sup> If this view is correct, then an effective altruist should count each abortion as an expected loss of over 70 quality-adjusted life years.<sup>33</sup>

Of course, we need to ask: *is* this view correct? We don't know. But suppose you are confidently pro-choice. You think that there is only a 10% chance that the pro-life view is correct.<sup>34</sup> What, then, is the expected social value of legalizing abortion?

Here's a back of the envelope calculation. Roughly 50 million abortions have been performed in the United States since *Roe v. Wade*.<sup>35</sup> Accounting for normative uncertainty, let's say that the expected cost of those abortions is 5 million lives lost. The U.S. Environmental Protection Agency estimates that the value of a statistical life is over \$7 million.<sup>36</sup> So the expected social cost of legalized abortion—even while being highly skeptical that a fetus's life should count morally—would be around \$35,000,000,000,000! (Of course, a full accounting would require that we also estimate the social benefits of legalized abortion—which further exacerbates the uncertainty.) If this is right, then a candidate's stance on abortion would override every other aspect of their platform. To be clear: we aren't sure that this approach is the right way to account for normative uncertainty—but uncertainty about uncertainty only compounds the problem.

We invoke this point not because we want to criticize or defend abortion. Our point is that the issue of normative uncertainty generalizes—there are many reasons to believe that determining the right way to vote is much harder than determining how to donate \$3000.

### *Cognitive Bias*

Even setting aside the problem of normative uncertainty, we have reason to doubt our capacities for accurately assessing which policies are welfare enhancing. Most of us are susceptible to *politically motivated reasoning*: we form political beliefs, not to arrive at the truth, but to protect and express our partisan identity.<sup>37</sup> Our vulnerability to politically motivated reasoning should drastically reduce our confidence in our beliefs about policy effectiveness.

Of particular interest here are findings indicating that our assessments of the effects and risks of policies are distorted by our partisan biases.<sup>38</sup> Proponents of capital punishment are resistant to evidence suggesting that it fails as a deterrent (the same holds for opponents, just reversed); proponents of handguns are resistant to evidence that gun control reduces crime (again, the same holds for opponents, just reversed).<sup>39</sup> Citizens of all political stripes are vulnerable to politically motivated reasoning.<sup>40</sup>

Critically, intelligence (including mathematical aptitude) appears to *increase* one's vulnerability to politically motivated reasoning.<sup>41</sup> People with the highest levels of political knowledge are especially likely to process political information in a biased way.<sup>42</sup> (Ignorant people engage in far less motivated reasoning than informed people.) These results may seem counterintuitive, but those with high intelligence and knowledge are better equipped to rationalize away counterevidence to their preferred political beliefs. Thus, it's not only unclear whether voting is a high expected value activity for the average citizen, but even for the sorts of people who read *Doing Good Better* or the *Journal of Political Philosophy*. The typical citizen in a "close enough" election should worry they lack the information needed to vote well, period, while the typical cognitively elite citizen should worry that they process that information in highly biased and unreliable ways.

Perhaps we are overthinking things. After all, remember that MacAskill suggests that we use "guesstimation" to assess the welfare impacts of different presidential administrations.<sup>43</sup> But this is a surprising and untenable move. MacAskill himself recognizes that people are bad at intuiting the effectiveness of policies. For instance, MacAskill's own research finds that people *think* that Scared Straight program works well, though it does great harm.<sup>44</sup> He found that Americans

answer barely better than chance when quizzed on the effectiveness of 10 government programs.<sup>45</sup> It is inconsistent, then, to show that people cannot properly intuit the effects of ten well-defined, existing programs but then suggest they can guesstimate the effects of *tens of thousands* of programs or decisions that a new presidential administration may undertake over four years.

Moreover, our “guesstimates” of policy outcomes are distorted not only by partisan biases but also by general cognitive biases. To take just one example, Cass Sunstein argues that “intuitive cost-benefit analysis” is impaired by the availability heuristic.<sup>46</sup> Vivid, readily-available information can exert an outsized influence on our judgment of a policies’ costs and benefits. For instance, seeing the body of someone killed in a nuclear meltdown overwhelms any statistical data one might have about the safety of nuclear power. Or suppose an administration implements an economic policy that boosts growth but also implements a military policy that gruesomely kills a few civilians. The latter might feature more prominently in our intuitive judgment even though the former matters more from a cost-benefit perspective.

You might be tempted by the thought that the preceding arguments miss the mark simply because one party is *obviously* better than the other. We’re skeptical that you should trust this thought. For one, the bias literature reveals that people exhibit unjustified self-confidence in their judgment in general and about politics in particular. Consider, for instance, that people are prone to *illusory superiority*, whereby we overestimate our abilities relative to others.<sup>47</sup> Moreover, the “bias blind spot” bias makes us believe that we are exempt from the biases that afflict others.<sup>48</sup> And learning about bias doesn’t reduce its influence.<sup>49</sup> Or take the *overconfidence effect*: our subjective confidence in our judgments tends to be higher than their actual accuracy.<sup>50</sup>

Further, research shows that people who score low in political knowledge are excessively confident in their political knowledge.<sup>51</sup> People generally think they understand the workings of public policies better than they actually do.<sup>52</sup> Finally, as we discussed above, rigorous empirical work assessing state-level differences in Democratic and Republican administrations fails to find much difference.

Diana Mutz finds that partisanship strongly predicts participation. If a citizen is correctly able to articulate and explain the other side's view, this predicts that the informed citizen rarely votes and contributes little to politics. If a citizen thinks instead that everyone on the other side is stupid and selfish, this predicts the citizen votes often and participates heavily.<sup>53</sup> Her data suggest that if one participates heavily in politics, one is probably overconfident.

Moreover, we believe that politically motivated reasoning supplies the best explanation for the thought that one party is clearly superior to the other. Consider the phenomenon of fact polarization. Dan Kahan writes,

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 [ . . . ] That explanation is *politically motivated reasoning*.<sup>54</sup>

The evidence that pertains to the effects of human activity on the climate is different than the evidence that pertains to the effects of gun ownership on crime. So why do these beliefs tend to go together?

The answer is that these beliefs have become signifiers of belonging to a particular political side. Kahan continues,

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.<sup>55</sup>

Our partisan identities determine how we interpret evidence—in brief, we accept evidence that affirms the correctness of our side and dismiss evidence that challenges the correctness of our side. Thus, it will *appear* as though all of the (credible) evidence supports our side regardless of whether it actually does. The best explanation for why you think your party is obviously better is not that it's obviously better, but that you are human.

#### **4. What is the Opportunity Cost of Voting Well?**

If votes matter, then voting is a risky activity, as any vote could be like giving or stealing thousands of dollars from charity. As we saw above, voters are beset by a host of cognitive and partisan biases, plus general problems of political ignorance, which impede their ability to judge whether they are helping or hurting. The empirical literature on bias even predicts that to claim “Sure, but I know the right answer, unlike most others” or “It's easy to judge which candidate is better” is unexceptional and a marker of cognitive bias and unreliability.

One might think, then, that the solution is for voters to gather more information and debias themselves. But these activities are costly. MacAskill and EGK claim that voting only costs an hour of your time. This is what it takes to register, travel to the polls, vote, and return home. However, the opportunity cost of voting *well* is much higher than an hour. Remember: if a vote this way is like donating \$3000 to a charity, a vote the other way is like destroying \$3000 of their funds.

It's worth asking how much information the typical American needs to get up to speed on politics. The short answer is *a lot*. Larry Bartels writes, “The political ignorance of the American

voter is one of the best-documented features of contemporary politics.”<sup>56</sup> Or, as political scientist Robert Luskin put the point, most Americans “know jaw-droppingly little about politics.”<sup>57</sup>

70 years of research on voter knowledge shows that the overwhelming majority of voters, including swing state voters in high stakes elections, lack any such knowledge. These facts are well-known, so we will not belabor the point here. Generally, voters in the US know who the president is and seem to remember and act upon about six months’ worth of the country’s economic performance (without knowing whether any particular incumbent was responsible for it), but little else. They do not know who their representatives are, what they did, what the effects of what they did were, what they propose to do, what they are likely to do, or much else.<sup>58</sup> One-third of Americans know next to nothing about politics; the typical American can answer basic political questions only slightly better than a random guesser, though a large percentage do significantly worse than chance.<sup>59</sup> Further, lay people’s beliefs about economics, political science, and natural science diverge significantly from those of experts,<sup>60</sup> and even experts are bad at making predictions.<sup>61</sup>

If you’re anything like the typical American, then, you might need to complete some online courses on economics and value theory to sharpen your competence as an evaluator of policies. This would take roughly 30 hours, if you understand everything on the first try. And suppose that you spend an hour per week tracking political news. That’s 52 hours. Still, more work remains. Acquiring political information does little good if we engage in politically motivated reasoning rather than arriving at true or justified conclusions. So to be a good voter, you need to invest time in debiasing. But the “easy” debiasing techniques—e.g., considering the opposite side—don’t work well and may actually backfire.<sup>62</sup> The most promising strategies are

time consuming. For instance, you can test your own understanding of the mechanics of different policies one by one and then systematically reevaluate them.<sup>63</sup> You might participate in highly curated formal conversations with out-group deliberately designed to depolarize participants.<sup>64</sup> But these sorts of techniques require considerable time and effort, raising their opportunity cost. Moreover, it remains unclear exactly how effective these techniques actually are. At most, they seem to increase epistemic humility, which is still just a start toward forming accurate political beliefs.

What, then, is the opportunity cost of becoming a good voter? For the sake of discussion, let's put the number at 100 hours. (We think that's extremely low for nearly all voters.) If you earn the mean American wage of \$26.00 and work 100 hours, that's \$2,600. Moreover, there is a good case to be made that effective altruists should invest their money to donate later rather than donating it immediately to maximize the good they can do over the course of their lifetime.<sup>65</sup> In that spirit, we will note that if you performed no further work and simply invested \$2,600 at 8% for 40 years, the total comes out to over \$56,000—which more than even the Charitable Voting Hypothesis generally claims you can bring about from swing state voting.

Lastly, as MacAskill notes, a good vote is, at most, like making a donation to a developed world charity. Thus, the beneficiaries of good votes are primarily other Americans who are wealthy by global standards (and most government domestic spending doesn't even target those who are poor by American standards). By contrast, the beneficiaries of effective giving are the severe poor. Because of the diminishing marginal utility of wealth, a dollar gain to a child in Ghana due to a donation to the Against Malaria Foundation is higher utility than a dollar gain to a US

citizen due to a good vote. Thus, we have another reason to believe that the expected utility of effective giving is greater than that of voting.

One might object that we're overestimating the costs of voting well precisely *because* the average voter is poorly informed—you don't need to learn much to become better than average.<sup>66</sup> In reply, we'll first note that it's simply unclear how much of a difference the information that is easy to acquire makes to forecasts of the welfare effects of a presidency when it isn't paired with information that is harder to acquire. For instance, it's not hard to learn what a candidate's stated position on charter schools is. But that information will only improve your vote if you also have knowledge about the costs and benefits of charter schools, knowledge which takes a greater investment to acquire. (We also worry that the other kinds of information you'd need to judge the utility difference between candidates are virtually impossible to acquire, such as those concerning counterfactual presidencies and knock-on effects of thousands of policies interacting in a complex world.)

Moreover, to increase the accuracy of one's political judgment, it's not enough to acquire information—one must also process that information in an unbiased way. Yet we've discussed evidence indicating that those citizens who are more likely to seek out political information are also more prone to bias. And the debiasing techniques available to us are costly to implement and of uncertain effectiveness. Of course, any given citizen may be a high-information, low-bias voter. But given how rare those are in general, you should be self-skeptical here, at least in the absence of strong evidence indicating that you are an exception to the rule.

*Heuristics*

Perhaps citizens need not learn economics and political philosophy themselves to evaluate the welfare effects of policies; they can simply use a heuristic.<sup>67</sup> For instance, they can follow the lead of a political party or a policy expert.

An initial problem with this reply is that people's selection of heuristics is swayed by their partisan political biases.<sup>68</sup> That progressives would use a high-profile progressive's opinion as a substitute for their own, but not a high-profile conservative's, is no surprise. Moreover, people's perception of who counts as an expert is similarly warped by politically motivated reasoning.<sup>69</sup> Of course, a citizen may always roll up their sleeves and do what's needed to arrive at an unbiased and well-informed judgment of which heuristics to trust. But now they've defeated the purpose of using a heuristic, which was to spare them the cost of doing their own epistemic work.<sup>70</sup>

Furthermore, even if a voter carefully selects an opinion leader that shares their values, using this heuristic need not lead them to vote for the higher utility candidate. For instance, a conservative voter might support capital punishment on retributivist grounds and defer to the opinion of a conservative columnist about the best retributivist candidate. However, the use of this heuristic wouldn't cause the voter to vote for the higher-utility candidate but rather the candidate that best satisfies their own retributivist preferences.

One solution to this problem is simply to develop something analogous to GiveWell.org for political candidates.<sup>71</sup> In this way, effective altruist voters could spare themselves the time of figuring out how to vote, just as GiveWell spares donors from figuring out how to donate. We have two replies. First, voting remains a low expected value activity until such a heuristic actually exists.

Second, the arguments we've made in the preceding sections cast doubt on the feasibility of developing such a heuristic. It is astronomically more difficult to assess potential presidential or parliamentary candidates than already operating charities. In further support of this claim, consider that there is a revealing feature of all existing effective altruist analyses of voting—they never specify which candidate is higher value. But this is exactly what voters would need. In the case of charitable giving, effective altruists tell us where to donate and show their work. But in the case of voting, they do neither.

To be clear, we recognize that organizations such as GiveWell are unable to determine the value of a donation with perfect mathematical precision.<sup>72</sup> They consider a wide range of factors and need to make judgment calls in their evaluations of charities. Yet we believe this point strengthens our case for the comparatively low value of voting. As complicated as it is to gauge the value of a donation to a single charity, it is *vastly* more complicated to gauge the value of a vote for a presidential candidate.

Consider another reason why raising funds for effective charities will tend to be higher utility than voting. The Charitable Voting Hypothesis itself rests on the assumption that, in a swing state, a single vote is high value. This assumption implies that a single vote for the Republican will be high expected instrumental value for the Republican but high expected instrumental *disvalue* for the Democrat. Thus, each swing-state vote for one side creates an incentive for the other party to expend additional resources mobilizing votes for their side.<sup>73</sup> As a result of this sort of arms race, casting even a meticulously researched vote may not make any difference because it will be neutralized by a countervailing effort from the outparty. Donating to effective charities is unlikely to give rise to a similar sort of arms race. After all, a donation to Deworm the

World won't prompt a neutralizing donation to Reworm the World (indeed, it's telling that this "anti-charity" doesn't even exist).

Lastly, if an election is very close, there is also a good chance that it will be decided by courts and lawsuits rather than the vote itself, which to some degree neutralizes the value of individual votes. In this case, an individual vote can make a difference to the outcome of the election by making the margin of victory sufficiently narrow, but it need not result in the election of the candidate preferred by the voter. Again, there is no equivalent in donations; it is not as though if a donation to Smile Savers would actually fix a cleft palate, then others will try to block the donation with lawsuits. Thus, there are considerations that further lower the expected value of a vote but which do not apply to charitable giving.

## **5. Conclusion**

*Outside* of politics, effective altruists generally say that we should recognize our poor epistemic states, recognize our biases, recognize our ignorance, and then direct our efforts toward cases where we are justified in believing we have a greater expected chance of doing real good. Accordingly, the new models of voter efficacy defended by MacAskill, Barnett, or EGK do not lead to triumphant conclusion that voters in swing states or close contests during high-stakes elections can regard themselves as donating the equivalent of \$3000 or \$50,000 to charity with a mere hour's work. Instead, they should see themselves as being in the dilemma of the biased, partly color-blind, highly unreliable, highly overconfident, and difficult-to-rehabilitate button-pushers who must choose between burgundy or maroon.

The good news about voting is that it matters more often, and matters more, than previously thought. The bad news is that it matters more often and matters more, which—given

the empirics of voter psychology and of assessing candidates—makes voting a high-risk, dangerous activity.

---

<sup>1</sup> For conventional accounts of the odds of casting a decisive vote, see e.g., Anthony Downs, *An Economic Theory of Democracy*. New York: Harper and Row, 1957; William Riker and Peter Ordeshook, “A Theory of the Calculus of Voting,” *American Political Science Review* 62 (1968): 25–42; John Ferejohn and Morris Fiorina, “The Paradox of Not Voting: A Decision Theoretic Analysis,” *American Political Science Review* 68 (1974): 525–536; Paul Meehl, “The Selfish Voter Paradox and the Thrown-Away Vote Argument,” *American Political Science Review* 71 (1977): 11–30; Guillermo Owen and Bernard Grofman, “To Vote or Not to Vote: The Paradox of Nonvoting,” *Public Choice* 42 (1984): 311–325; John Aldrich, “Rational Choice and Turnout,” *American Journal of Political Science* 37 (1993): 246–278; Donald Green and Ian Shapiro, *Pathologies of Rational Choice Theory*. New Haven: Yale University Press, 1994; Alan Carling, “The Paradox of Voting and the Theory of Social Evolution,” in *Preferences, Institutions and Rational Choice*, Keith Dowding and Desmond King, eds. Oxford: Clarendon Press, 1995. Pp. 20-43; Ron Shachar and Barry Nalebuff, “Follow The Leader: Theory and Evidence on Political Participation,” *American Economic Review* 89 (1999): 525–547; Dennis Mueller, *Public Choice III*. Cambridge: Cambridge University Press, 2003, 305.

<sup>2</sup> Aaron Edlin, Andrew Gelman, and Noah Kaplan, “Vote for Charity’s Sake,” *The Economists’ Voice* 5 (2008): 1-4, 4.

<sup>3</sup> William MacAskill, *Doing Good Better: How Effective Altruism Can Help You Make a Difference*. New York: Penguin, 2015, 86.

<sup>4</sup> Zach Barnett, “Why You Should Vote to Change the Outcome,” *Philosophy and Public Affairs* 48 (2020): 422-446.

<sup>5</sup> For another argument in this spirit, see Derek Parfit, *Reasons and Persons*. Oxford: Oxford University Press, 1984, 73-5.

<sup>6</sup> Jason Brennan, “The Ethics and Rationality of Voting,” *Stanford Encyclopedia of Philosophy*, 2016 edition, ed. Edward N. Zalta, URL = <https://plato.stanford.edu/entries/voting/#1>.

<sup>7</sup> Journalists sometimes claim that votes can also have value in helping to “change the mandate,” but empirical work shows that a politician’s margin of victory has no effect on her effectiveness in office. See Dahl 1990; Noel 2008. Anthony Downs (2006, 257) conjectured that perhaps casting a vote has some chance of stopping a democracy from collapsing, but we will follow Brennan (2011, 23-28), who argues that no matter how we interpret Downs’s conjecture, nevertheless, the expected utility of a vote in terms of its tendency to preserve democracy is vanishingly small. So, for this paper, we will put these issues aside.

<sup>8</sup> Geoffrey Brennan and Loren Lomasky, *Democracy and Decision* (Cambridge University Press, 1993).

<sup>9</sup> Barnett, 2020; Aaron Edlin, Andrew Gelman, and Noah Kaplan, “Voting as a Rational Choice: Why and How People Vote to Improve the Well-Being of Others,” *Rationality and Society* 19 (2007): 219-314.

- 
- <sup>10</sup> Edlin, et al 2007.
- <sup>11</sup> Jason Brennan, *The Ethics of Voting* (Princeton University Press, 2011), 12.
- <sup>12</sup> Loren Lomasky and Geoffrey Brennan, "Is There a Duty to Vote?", *Social Philosophy and Policy* 17 (2000): 62-86, 71.
- <sup>13</sup> MacAskill 2015, 86.
- <sup>14</sup> This example takes inspiration from Christopher Freiman, *Why It's OK to Ignore Politics*. New York: Routledge, 2020, chapter 3.
- <sup>15</sup> Edlin, Gelman, and Kaplan 2008, 2. Emphasis ours.
- <sup>16</sup> MacAskill 2015, 85.
- <sup>17</sup> On the Playpump, see MacAskill 2015, 1-5. On Scared Straight, see William MacAskill, Benjamin Todd, and Robert Wiblin, "Can You Guess Which Government Programs Work? Most People Can't," Vox. 8.17.15. Accessed 6.6.19. Available online at: <https://www.vox.com/2015/8/13/9148123/quiz-which-programs-work>; Steve Aos, Roxanne Lieb, Jim Mayfield, Marna Miller, and Annie Pennucci, *Benefits and Costs of Prevention and Early Intervention Programs for Youth*. Olympia: Washington State Institute for Public Policy, 2004.
- <sup>18</sup> MacAskill 2015, 85.
- <sup>19</sup> Philip Tetlock, *Expert Political Judgment*. Princeton: Princeton University Press, 2005.
- <sup>20</sup> Thanks to [omitted] for pointing out the significance of the existence of prediction markets.
- <sup>21</sup> 2000, 71.
- <sup>22</sup> Gerald Gaus, "What Might Democratic Self-Governance in a Complex Social World Look Like?" *San Diego Law Review* 56 (2019): 967-1012, 975.
- <sup>23</sup> Brian Burgoon, "Immigration, Integration, and Support for Redistribution in Europe," *World Politics* 66 (2014): 365-405; A. Alesina and E. La Ferrara, "Preferences for Redistribution in the Land of Opportunities," *Journal of Public Economics* (2005): 897-931; A. Alesina, E. Glaeser, and B. Sacerdote, "Why Doesn't the US Have a European-Style Welfare State?," *Brookings Papers on Economic Activity* 32 (2001): 187-277; Martin Gilens, "Racial Attitudes and Opposition to Welfare," *Journal of Politics* 57 (1995): 995-1014. For a discussion of this point in the context of complexity in particular, see Jacob Barrett, "Social Reform in a Complex World," *Journal of Ethics and Social Philosophy* 17 (2020): 103-132, 109.
- <sup>24</sup> E.g., Glen Weyl, "The Openness-Equality Trade-Off in Global Redistribution," *The Economic Journal* 128 (2018): F1-F36.
- <sup>25</sup> See, for instance, *The Economist*, "Explaining the Brexit Vote," 7.14.16. Available online at: <https://www.economist.com/britain/2016/07/14/explaining-the-brexit-vote>. Accessed 1.4.20.
- <sup>26</sup> For evidence that this is the case for some American voters, see Eli Yokley, "Ahead of 2020, Democratic Voters are Moving Left on Immigration," *Morning Consult* 8.28.19 Available online at: <https://morningconsult.com/2019/08/28/ahead-2020-democratic-voters-moving-left-immigration/>. Accessed 1.4.20.
- <sup>27</sup> 2020, 428.
- <sup>28</sup> Melvyn Leffler, "9/11 in Retrospect," *Foreign Affairs* 90 (2011): 33-44.
- <sup>29</sup> Robert Wiblin, "If You Care About Social Impact, Why is Voting Important?" *80,000 Hours*. 10/20. Available online at: <https://80000hours.org/articles/is-voting-important/>. Accessed 7/28/21.
- <sup>30</sup> Adam Dyne and John Holbein, "Noisy Retrospection: The Effect of Party Control on Policy Outcomes," *American Political Science Review* 114 (2020): 237-257, here 237. There are

---

important differences between state and national elections that produce a larger value difference between candidates in the latter. For example, it may be the case that governors simply have less of an effect on their constituents' welfare than presidents and that presidents have significantly more impact on the welfare of non-citizens as a result of foreign policy decisions. For these sorts of reasons, we should regard the state-level evidence as suggestive but far from dispositive. We are grateful to an anonymous referee for raising these concerns.

<sup>31</sup> We thank an anonymous referee for raising this point.

<sup>32</sup> Don Marquis, "Why Abortion is Immoral," *Journal of Philosophy* 86 (1989): 183-202.

<sup>33</sup> For a similar analysis of abortion from the perspective of effective altruism, see "Compartmentalizing: Effective Altruism and Abortion," *LessWrong*. 1.4.15. Available online at: <https://www.lesswrong.com/posts/E6dDvRCr8eeTDJrAG/compartmentalizing-effective-altruism-and-abortion>. Accessed 1.4.20.

<sup>34</sup> For more on abortion and moral risk, see Dan Moeller, "Abortion and Moral Risk," *Philosophy* 86 (2011): 425-443.

<sup>35</sup> Data from Rachel Jones and Jenna Jerman, "Abortion Incidence and Service Availability In the United States, 2014," *Perspectives on Sexual and Reproductive Health* 49 (2017): 1-11. Available online at: [https://www.guttmacher.org/sites/default/files/article\\_files/abortion-incidence-us.pdf](https://www.guttmacher.org/sites/default/files/article_files/abortion-incidence-us.pdf)

<sup>36</sup> United States Environmental Protection Agency, "Mortality Risk Valuation," Available online at: <https://www.epa.gov/environmental-economics/mortality-risk-valuation#whatvalue>. Accessed 1.4.20.

<sup>37</sup> See Charles Lord, Lee Ross, and Mark Lepper, "Biased Assimilation and Attitude Polarization: Effects of Prior Theories on Subsequently Considered Evidence," *Journal of Personality and Social Psychology* 37 (1979): 2098-2109; Larry Bartels, "Beyond the Running Tally: Partisan Bias in Political Perceptions," *Political Behavior* 24 (2002): 117-50; David Redlawsk, "Hot Cognition or Cool Consideration?" *Journal of Politics* 64 (2002): 1021-1044; Geoffrey Cohen, "Party Over Policy: The Dominating Impact of Group Influence on Political Beliefs," *Journal of Personality and Social Psychology* 85 (2003): 808-822; Charles Taber and Milton Lodge, "Motivated Skepticism in the Evaluation of Political Beliefs," *American Journal of Political Science* 50 (2006): 755-769; David Sherman and Geoffrey Cohen, "The Psychology of Self-Defense: Self-Affirmation Theory," in Mark Zanna, ed. *Advances in Experimental Social Psychology*, vol. 38. San Diego, CA: Academic Press, 2006, 183-242; Alan Gerber and Gregory Huber, "Partisanship, Political Control, and Economic Assessments." *American Journal of Political Science* 54 (2010): 153-73; Brendan Nyhan and Jason Reifler, "When Corrections Fail: The Persistence of Political Misperceptions," *Political Behavior* 32 (2010): 303-330; Dan Kahan, "Ideology, Motivated Reasoning, and Cognitive Reflection," *Judgment and Decision Making* 8 (2013): 407-424; Dan Kahan, "Climate-Science Communication and the Measurement Problem," *Political Psychology* 36 (2015) 1-43; Dan Kahan, "The Politically Motivated Reasoning Paradigm, Part 1: What Politically Motivated Reasoning Is and How to Measure It," *Emerging Trends in the Social and Behavioral Sciences*, eds R. A. Scott and S. M. Kosslyn, 2016a doi:10.1002/9781118900772.etrds0417; Dan Kahan, "The Politically Motivated Reasoning Paradigm, Part 2: Unanswered Questions," In *Emerging Trends in the Social and Behavioral Sciences*, eds R. A. Scott and S. M. Kosslyn. 2016b. doi:10.1002/9781118900772.etrds0418; Dan Kahan, Donald Braman, Geoffrey Cohen, John Gastil, Paul Slovic, "Who Fears the HPV Vaccine, Who Doesn't, and Why? An Experimental Study

---

of the Mechanisms of Cultural Cognition,” *Law and Human Behavior* 34 (2010): 501-516; Dan Kahan, Hank Jenkins-Smith, and Donald Braman, “Cultural Cognition of Scientific Consensus,” *Journal of Risk Research* 14 (2011): 147-174; Dan Kahan, David Hoffman, Donald Braman, Danieli Evans, and Jeffrey Rachlinski, “They Saw A Protest: Cognitive Illiberalism and the Speech-Conduct Distinction,” *Stanford Law Review* 64 (2012): 851-906; Dan Kahan, Ellen Peters, Erica Cantrell Dawson, and Paul Slovic, “Motivated Numeracy and Enlightened Self-Government,” *Behavioural Public Policy* 1 (2017): 54-86; Jacob Westfall, Leaf Van Boven, John Chambers, and Charles Judd, “Perceiving Political Polarization in the United States Party Identity Strength and Attitude Extremity Exacerbate the Perceived Partisan Divide,” *Perspectives on Psychological Science* 10 (2015): 145–158; Christopher Achen and Larry Bartels, *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton: Princeton University Press, 2016; Jay Van Bavel and Andrea Pereira, “The Partisan Brain: An Identity-Based Model of Political Belief,” *Trends in Cognitive Sciences* 22 (2018): 213-224.

<sup>38</sup> See Lord et al, 1979; Kahan et al, 2011.

<sup>39</sup> See Lord et al, 1979; Kahan et al, 2017.

<sup>40</sup> Kahan, 2013.

<sup>41</sup> Ibid.

<sup>42</sup> Taber and Lodge 2006.

<sup>43</sup> MacAskill 2015, 85.

<sup>44</sup> Aos et al, 2004, 7.

<sup>45</sup> William MacAskill, Benjamin Todd, and Robert Wiblin, “Can You Guess Which Government Programs Work? Most People Can’t,” *Vox*. 8.17.15. Accessed 6.6.19. Available online at: <https://www.vox.com/2015/8/13/9148123/quiz-which-programs-work>

<sup>46</sup> Cass Sunstein, “The Availability Heuristic, Intuitive Cost-Benefit Analysis, and Climate Change,” *Climatic Change* 77 (2006): 195–210.

<sup>47</sup> See, e.g., Vera Hoorens, “Self-Favoring Biases, Self-Presentation, and the Self-Other Asymmetry in Social Comparison,” *Journal of Personality* 63 (1995): 793-817.

<sup>48</sup> Emily Pronin, Daniel Lin, and Lee Ross, “The Bias Blind Spot: Perceptions of Bias in Self Versus Others,” *Personality and Social Psychology Bulletin* 28 (2002): 369-381.

<sup>49</sup> Ibid.

<sup>50</sup> Gerry Pallier, Rebecca Wilkinson, Vanessa Danthiir, Sabina Kleitman, Goran Knezevic, Lazar Stankov, and Richard Roberts, “The Role of Individual Differences in the Accuracy of Confidence Judgments,” *The Journal of General Psychology* 129 (2002): 257–299.

<sup>51</sup> Ian Anson, “Partisanship, Political Knowledge, and the Dunning-Kruger Effect,” *Political Psychology* 39 (2018): 1173-1192.

<sup>52</sup> Fernbach et al, 2013.

<sup>53</sup> Diana Mutz, *Hearing the Other Side* (Cambridge University Press, 2006).

<sup>54</sup> Kahan 2016a, 2. Emphasis in the original.

<sup>55</sup> Ibid. Kahan cites Sherman and Cohen 2006.

<sup>56</sup> Larry Bartels, “Uninformed Votes: Information Effects in Presidential Elections.” *American Journal of Political Science* 40 (1996): 194-230, 194.

<sup>57</sup> Robert Luskin, “From Denial to Extenuation (and Finally Beyond): Political Sophistication and Citizen Performance.” In James Kuklinski, ed. *Thinking About Political Psychology*. New York: Cambridge University Press, 2002, 281-305, 282.

---

<sup>58</sup> Michael X. Delli Carpini and Scott Keeter, *What Americans Know about Politics and Why It Matters* (Yale University Press, 1997); Ilya Somin, *Democracy and Political Ignorance* (Stanford University Press, 2013).

<sup>59</sup> Somin, 2013, 33.

<sup>60</sup> See, respectively, Bryan Caplan, *The Myth of the Rational Voter*. Princeton: Princeton University Press, 2007, chapter 3; Bryan Caplan, Eric Crampton, Wayne Grove, and Ilya Somin, "Systematically Biased Beliefs about Political Influence: Evidence from the Perceptions of Political Influence on Policy Outcomes Survey," *PS: Political Science & Politics* 46 (2013): 760-767.; Pew Research Center, "Elaborating on the Views of AAAS Scientists, Issue by Issue," 7.23.15. Accessed 6.25.19. Available online at: <https://www.pewresearch.org/science/2015/07/23/elaborating-on-the-views-of-aaas-scientists-issue-by-issue/>

<sup>61</sup> Tetlock 2005.

<sup>62</sup> Redlawsk 2002; see also Michael Meffert, Sungeun Chung, Amber Joiner, Leah Waks, and Jennifer Garst, "The Effects of Negativity and Motivated Information Processing During a Political Campaign." *Journal of Communication* 56 (2006): 27–51; Christopher Bail, Lisa Argyle, Taylor Brown, John Bumpus, Haohan Chen, M. B. Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky, "Exposure to Opposing Views on Social Media Can Increase Political Polarization," *Proceedings of the National Academy of Sciences* 115 (2018): 9216-9221.

<sup>63</sup> Philip Fernbach, Todd Rogers, Craig Fox, Steven Sloman, "Political Extremism Is Supported by an Illusion of Understanding," *Psychological Science* 24 (2013): 939-946.

<sup>64</sup> Peter Coleman, *The Five Percent: Finding Solutions to Seemingly Impossible Conflicts*. New York: PublicAffairs 2011.

<sup>65</sup> Dan Moeller, "Should We Let People Starve – For Now?" *Analysis* 66 (2006): 240-247.

<sup>66</sup> We are grateful to an anonymous referee for making this objection.

<sup>67</sup> Arthur Lupia and Matthew McCubbins. *The Democratic Dilemma: Can Citizens Learn What They Need To Know?* Cambridge: Cambridge University Press, 1998; Thomas Christiano, "Voter Ignorance Is Not Necessarily a Problem." *Critical Review* 27 (2015): 253-269, 257-260.

<sup>68</sup> Ilya Somin, *Democracy and Political Ignorance*. Stanford: Stanford University Press, 2013, 99-100.

<sup>69</sup> Kahan et al, 2011.

<sup>70</sup> Somin 2013, 99.

<sup>71</sup> Thanks to an anonymous referee for making this suggestion.

<sup>72</sup> We are grateful to an anonymous referee for emphasizing this point.

<sup>73</sup> Thanks to [omitted] for raising this point.