

ZACH BARNETT* 

Why You Should Vote to Change the Outcome

I. INTRODUCTION

When voting comes with a cost, why pay it? Sometimes, there is a simple answer. We pay the cost *to make our preferred outcomes likelier*. When I cast my vote for Class President or Team Captain, there's a certain result I want, and I'm trying to bring it about.

One might worry, though, that this simple rationale is inapplicable to very large elections. After all, when there are millions and millions of voters, the chance that my individual vote will make a difference to the final outcome is utterly miniscule. How could it be rational for me to do something that's virtually certain to have zero impact?

In response to this challenge, one might encourage me not to overlook the magnitude of the *stakes*. If there are millions of voters, there are, presumably, millions of people who will be affected by the result. Yes, the chance that my vote makes the difference is very tiny, but the difference my vote could make is very great. Arguably, the magnitude of the stakes can, at least sometimes, offset the tininess of the chance of affecting the

Zach Barnett is an Assistant Professor in the philosophy department at the National University of Singapore, where he teaches ethics, epistemology, and decision theory. For helpful comments and suggestions, I'd like to acknowledge Nomy Arpaly, Bob Beddor, Ben Blumson, Jason Brennan, Anna Brinkerhoff, David Builes, Nevin Climenhaga, Jamie Dreier, Stephanie van Fossen, Simon Goldstein, Preston Greene, John Hawthorne, Alan Hajek, Brian Hedden, Qu Hsueh, Ethan Jerzak, Han Li, Jowen Lim, Chad Marxen, Daniel Muñoz, Julia Nefsky, Mike Pelczar, Abelard Podgorski, Gerard Rothfus, Josh Schechter, Neil Sinhababu, Louanne Sungaimin, Teoh Tai Ker, Margot Witte, Patrick Wu, the Faculty Reading Group at NUS, audiences at Nanyang Technological University, the Probability in Philosophy Workshop at Australian Catholic University, and the editors of *Philosophy & Public Affairs*. Special thanks to David Christensen and Dave Estlund.

outcome—making it rational to vote solely in virtue of the expected consequences of doing so. This is the *consequentialist defense* of voting.¹

While it would certainly be nice if the consequentialist defense could succeed, the prevailing view seems to be that, unfortunately, the numbers just don't add up.² Here is a representative passage from the *Stanford Encyclopedia of Philosophy* entry on "The Ethics and Rationality of Voting," authored by Jason Brennan.

There is some debate among economists and political scientists over the precise way to calculate the probability that a vote will be decisive. Nevertheless, they generally agree that the probability that the modal individual voter in a typical election will break a tie is small, so small that the expected benefit . . . of the modal vote for a good candidate is worth far less than a millionth of a penny. [. . .] Thus . . . for most voters in most elections, voting for the purpose of trying to change the outcome is irrational. The expected costs exceed the expected benefits by many orders of magnitude.³

Elsewhere, Brennan offers a compelling example, purportedly illustrating the futility of voting to change the outcome—even when the stakes are assumed to be extraordinarily high.⁴ Brennan's example is succinctly summarized by Luke Maring:

1. Parfit sketches a defense of voting along these lines in *Reasons and Persons* (Oxford: Oxford University Press, 1984), 73–4. But today, few in philosophy endorse it.

2. Alexander Guerrero deems the consequentialist defense "unsatisfactory" ("The Paradox of Voting and the Ethics of Political Representation," *Philosophy & Public Affairs* 38 (2010): 274); Luke Maring considers it "not promising" ("Debate: Why Does the Excellent Citizen Vote?" *Journal of Political Philosophy* 24 (2016): 245); Geoffrey Brennan and Geoffrey Sayre-McCord assert that the prospect of casting the deciding vote will never move the rational voter ("Voting and Causal Responsibility," *Oxford Studies in Political Philosophy* 1 (2016): 38); Christopher Freiman argues that one's duty to vote "cannot be based on the good political consequences" of one's doing so (*Why It's OK to Ignore Politics* (New York: Routledge, 2020), chapter 3).

3. There are some departures from this consensus, such as 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, who make an empirical case for the rationality of voting. In Section IV, we will examine the research to which Brennan alludes, as it relates to the expected impact of a single vote.

4. Jason Brennan, *The Ethics of Voting* (Princeton: Princeton University Press, 2011), 18–20.

[I]magine that a particular candidate's victory is worth \$33 billion to the common good (suppose she is a civic-minded, financial wizard), that there are 122,293,322 voters (as in the 2004 U.S. presidential election), and that the probability of any given voter supporting our financial wizard is 50.5%. With the stakes artificially raised, one might expect that individual votes are impactful. But the expected value to the common good of one's vote for the financial wizard is a mere $\$4.77 \times 10^{-2650}$. We might wonder whether expected financial return is the best way to measure the value of casting a ballot. But however those wonderings turn out, Brennan's example illustrates that an individual vote is a drop in the ocean.⁵

To drive the point home, Brennan observes that while the nucleus of an atom is about fifteen orders of magnitude smaller than a human being, the expected benefit, in dollars, of a vote for the financial wizard is 2,648 orders of magnitude smaller than a penny. If we are looking to justify voting, it is tempting to conclude that we've no choice but to look elsewhere.⁶

Such a conclusion would be premature, however. A simple yet powerful consequentialist case for voting can be made. Given certain basic assumptions, the rationality of voting can be proven to hold given two conditions: a *stakes* condition and a *chances* condition. It will be argued here that both conditions are often satisfied in typical electoral circumstances. After examining the argument, we will conclude with a discussion of why Brennan's example led us astray.

II. THE EXPECTED VALUE OF VOTING

Suppose you're eligible to vote in a large, upcoming election. On the ballot are two candidates: Daisy and Donald. You regard Daisy's policies as

5. Maring, "Why Does the Excellent Citizen Vote?" 245.

6. Owing partly to widespread pessimism about the prospects for a consequentialist defense of voting, some have sought to defend voting in other ways. Specifically, some have suggested that we vote to express our preferences (Geoffrey Brennan and Loren Lomasky, *Democracy and Decision: The Pure Theory of Electoral Preference* (Cambridge: Cambridge University Press, 1993)), to share in a winning effort (Alvin Goldman, "Why Citizens Should Vote: A Causal Responsibility Account," *Social Philosophy and Policy* 16 (2001): 201-17 and Richard Tuck, *Free Riding* (Cambridge: Harvard University Press, 2008)), to be an excellent citizen (Richard Dagger, *Civic Virtues: Rights, Citizenship, and Republican Liberalism* (Oxford: Oxford University Press, 1997)), or to respect democracy (Maring, "Why Does the Excellent Citizen Vote?").

significantly better for the common good. You'll vote for her, provided that you vote at all. But you're not sure whether you'll vote. You're inclined to do so only to the extent that your vote bears directly on the final result. And in all likelihood, it won't. At the same time, you recognize that the stakes are high. Should you vote? How should you decide? In thinking through your predicament, we will make two simplifying assumptions, which will inform our discussion.

To make the problem tractable, we will make the standard decision-theoretic assumption that *an act is rational if its expected benefits exceed its costs*. But we must specify: costs and benefits to whom—you, the voter? No. If you were wholly self-interested, voting would hardly ever be rational.⁷ Since the aim of this article is to assess the consequentialist case for voting, we will assume that *you are a consequentialist voter*—that is, your voting decisions (whether to vote, whom to vote for) will be based solely on what is best for the public, where everyone's interests are given equal weight, including your own.⁸ So when we're assessing the expected impact of your potential vote, we'll be thinking about the social costs and the social benefits.

To determine whether you should vote or abstain, we'll need to compare the cost of your voting, whatever it is, with the expected social benefit of a vote for Daisy. If this expected benefit exceeds the cost, your voting for Daisy is rational.⁹

7. This fact gives rise to Anthony Downs' *paradox of voting*, sometimes called the *paradox of voter turnout*. Downs, *An Economic Theory of Democracy* (New York: Harper & Brothers, 1957). On the classical assumptions of self-interest and rationality, the fact that anyone votes at all is difficult to explain. Our focus, however, is on a different issue—whether voting is rational for a *consequentialist voter*, as specified below.

8. I'm working, here, under a means-end conception of rationality, according to which potential courses of action are rational to the extent that they promote the agent's aims or goals, whatever those happen to be. In the case at hand, the thought is that some citizens may be public-spirited voters: they may aim to benefit the public through their voting decisions. For such people, the rationality of voting will depend upon whether the act of voting actually tends to benefit the public. Brennan and others argue that, even in the best case, the cost imposed upon the voter exceeds (in expectation) the benefit conferred to the rest of the public; this article defends an alternative view.

9. Strictly speaking, for voting to be rational, it must have a higher expected value than all possible alternative courses of action. For simplicity, the only options we will consider are *voting* and *abstaining* (which are, I think, the two main options on many would-be voters' minds). If it could be shown that, rationally speaking, the public-spirited voter should vote rather than stay at home, this would be a noteworthy observation—and one which runs contrary to received wisdom.

The expected benefit of a vote for Daisy is basically a function of two variables. First, there's the *net social benefit* associated with having Daisy rather than Donald in charge, which we can call B . Second, there's the *probability that your vote for Daisy is decisive*, which we can call d .

What does it mean for a vote to be decisive? We'll follow Brennan and others in using the term "decisive" to refer to a vote that *changes the outcome*—that is, a vote that breaks a tie (when the number of voters is odd) or causes a tie (when the number of voters is even).

Now it might seem that the expected benefit of a vote for Daisy, then, would simply be the product of our two variables, $B \times d$. But given how we have understood decisiveness, it's actually going to be about half of that.

Why? Suppose your vote turns out to be decisive: it breaks a tie in favor of Daisy. Even here, your vote doesn't cause Daisy to win *when she otherwise would have lost*. If you had abstained, there would have been a tie. What happens in case of ties? For simplicity, we'll assume that, in the event of an exact tie, both candidates are equally likely to be awarded victory—perhaps a coin is flipped to determine a winner. Given this assumption, the expected value of a tie between Daisy and Donald is $\frac{1}{2} \times B$. It follows that the expected benefit of a decisive vote for Daisy, one which causes Daisy to win *when she otherwise would have tied*, is $B - \frac{1}{2} \times B$, or simply $\frac{1}{2} \times B$.¹⁰

So the expected benefit of a vote for Daisy, then, is $\frac{1}{2} \times B \times d$. For voting to be rational, this expected benefit must exceed the voting cost, c .

$$EB_{\text{vote}} = \frac{1}{2} \times B \times d > c.$$

In thinking through how to assess whether this condition is met in a given case, it will prove helpful to think about how the expected benefit of voting varies with N , the number of citizens in the community.

The probability that your vote is decisive, d , diminishes as N grows. The more voters there are, the smaller your chance of being the difference-maker.

In contrast, the net social benefit, B , is proportional to N . After all, an election that affects eight million is, other things equal, twice as important to the consequentialist voter, as an election that affects four million. Accordingly, we will rewrite B as $b \times N$, where b stands for the *average*

10. When N is even, a decisive vote causes Daisy to tie when she otherwise would have lost. Here too, the expected benefit of a decisive vote works out to $\frac{1}{2} \times B$.

social benefit—that is, the average benefit, per person—of having Daisy rather than Donald in charge.

With this in mind, here are two substantive conditions, which happen to be together sufficient for the rationality of voting.

First, there is the *stakes condition*, which requires the average social benefit of electing the better candidate to be more than twice as great as the individual voting cost.

Stakes condition: The average social benefit of electing the better candidate is more than twice as great as the cost of voting (in short: $b > 2 \times c$).

In Section III, it will be argued that this condition is often satisfied in the real world.

Second, there is the *chances condition*, which requires the chance of casting a decisive vote to be at least one divided by the number of citizens.

Chances condition: The probability of casting the deciding vote is at least one divided by the number of citizens (in short: $d \geq 1/N$).

On the face of it, this condition may not seem plausible for large elections. But in Section IV, it will be argued that this condition, too, is often met in the real world.

Given these two conditions, the expected benefit of voting can be proven to be greater than the cost.

$$EB_{\text{vote}} = \frac{1}{2} \times B \times d = \frac{1}{2} \times b \times N \times d \geq \frac{1}{2} \times b \times 1 = \frac{1}{2}b > \frac{1}{2}c$$

(inequality *i* follows from the chances condition; inequality *ii* follows from the stakes condition.)

The bottom line: if the stakes and chances conditions are met, voting is rational. How often are these conditions true in the real world? We'll explore that next.

III. THE STAKES CONDITION: A QUALIFIED DEFENSE

The stakes condition asserts that the average social benefit of electing the better candidate is more than twice as great as the individual cost of

voting. How often is this true in real life? Not always, to be sure. There isn't always very much at stake. For instance, two candidates' respective platforms may be quite similar. Or even if they differ, the differences may wash out. If, for example, Daisy's policies stand to benefit rural voters, while Donald's stand to benefit urban voters, the net social benefit associated with electing Daisy might be slim to none.

Nevertheless, there is reason to think that the stakes condition will be met in a wide range of realistic electoral circumstances. Consider, for example, a simple referendum on a policy substantially reducing the tax burden on every household earning less than the median income. Though such a policy would not benefit all citizens, it would benefit about half of them quite substantially. The average benefit of this tax relief policy (provided that the downsides were small) would be quite high—plausibly much higher than the typical cost of voting.

Or consider a different example. Estimates of the total cost of military operations in Iraq and Afghanistan undertaken by the American Government in the aftermath of September 11, 2001 range from \$1.4 trillion to \$3.5 trillion.¹¹ This is more than \$4000 per U.S. citizen. Even if these wars delivered significant benefits to the public (a proposition which would be emphatically denied by some), the sheer magnitude of these figures illustrates just how much can ride on political decisions. The individual cost of voting pales in comparison.

Here's still another way to think about the question. According to the Office of Management and Budget, the U.S. Government generates more than three trillion dollars per year in tax revenue—about \$10,000 per U.S. citizen.¹² The result of an American presidential election affects just how this revenue will be spent. If one candidate's spending profile were,

11. The lower estimate is offered by the Congressional Budget Office of the United States, "Estimated Costs of U.S. Operations in Iraq and Afghanistan and of Other Activities Related to the War on Terrorism," statement of Peter Orszag, testimony before the Committee on the Budget of the United States House of Representatives (2007); the higher one is offered by the Joint Economic Committee of the United States Congress, "War at Any Price? The Total Economic Cost of the War Beyond the Federal Budget," *Joint Economic Committee of the United States Congress*, published November 2007, <https://www.hsdl.org/?abstract&did=480904/>. See also Joseph Stiglitz and Linda Bilmes, *The Three Trillion Dollar War* (New York: W.W. Norton & Company, 2008), who price the Iraq War at \$3 trillion.

12. Office of Management and Budget of the United States, *Budget of the United States Government, Fiscal Year 2019* (Washington, DC: United States Government Publishing Office, 2018), 90.

say, 5 percent more efficient than another's, it would amount to a difference of \$500/citizen in well spent revenue. Presumably, an average social benefit of \$500 would dwarf the individual cost of voting for most Americans.

The bottom line: in at least some fairly ordinary electoral circumstances, the stakes condition is comfortably met.¹³

IV. THE CHANCES CONDITION: HOW LIKELY AM I TO CAST A DECISIVE VOTE?¹⁴

At this point, everything seems to be riding on the chances condition. Is it really true that one's chance of casting a decisive vote is greater than $1/N$?

Often, the answer is "yes." This may seem surprising. But in this section, we'll see that the chances condition follows from two weak modeling assumptions, together with the claim that *both candidates have at least a 10 percent chance of winning*. Before presenting the argument, it is advisable to examine other ways of estimating d , the chance of casting a decisive vote.

13. Even if it is conceded that the *actual* social benefits of electing the better candidate are sufficiently large, one might object that the stakes condition is still not necessarily met on the grounds that voters may not be in a position to *foresee* these benefits. Arguments along these lines have been developed by Bryan Caplan, *The Myth of the Rational Voter: Why Democracies Choose Bad Policies* (Princeton, NJ: Princeton University Press, 2008); Brennan, *The Ethics of Voting*; Brennan, *Against Democracy* (Princeton, NJ: Princeton University Press, 2016); Michael Huemer, "In Praise of Passivity," *Studia Humana* 1 (2012): 12–28; Freiman, *Why It's OK to Ignore Politics*. These authors make the case against voting from the premise of voter ignorance. At first glance, their conclusions seem in direct conflict with the view defended in this article. But upon closer examination, the two positions may actually be compatible. When we ask whether voting is rational, we should distinguish two questions: (1) Is it rational to vote, *given voters' actual beliefs*? (2) Is it rational to vote, *given the beliefs voters ought, epistemically, to have*? To see why this distinction matters, suppose that Charlie foolishly believes that Daisy is the messiah and that he has a 50 percent chance of single-handedly tipping the election by voting for her. Holding fixed Charlie's irrational beliefs, of course voting has positive expected value. But this facile response leaves an important question unanswered: it does not tell us whether voting is rational, given a *realistic* assessment of the candidates and a *realistic* estimate of one's chance of casting a decisive vote. This suggests that question (2) is important. To the extent that we are interested in (2), we are entitled to leave aside worries about voters' actual beliefs and simply assume that our imagined voter is a well-informed citizen.

14. Thanks to the editors of *Philosophy & Public Affairs* whose constructive suggestions clarified and enriched the arguments advanced in this section.

A. Existing Approaches

Some of the earliest estimates of d are found in discussions of *voting power*—a measure of how much control a given citizen has over electoral outcomes.¹⁵ For example, in the United States, we might wonder whether residents of different states possess equal voting power in presidential elections.¹⁶ One popular way to investigate questions of this sort involves supposing that each citizen votes at random. That is, each voter's decision is modeled as an independent coin toss, with a 50 percent chance of producing a vote for either candidate.¹⁷ This model, sometimes called *random voting* (or *impartial culture*, among social choice theorists¹⁸), can be used to estimate the value of d . If random voting is assumed, d actually turns out to be far greater than $1/N$. For example, if you and 500,000 fellow citizens all vote randomly, your chance of casting a decisive vote is about 1 in 1,250.¹⁹ In general, under random voting, d is proportional to $1/\sqrt{N}$.²⁰

But random voting, as described above, is not a flexible model. In effect, it assumes that both candidates have *exactly* the same chance of

15. Lionel Penrose, "The Elementary Statistics of Majority Voting," *Journal of the Royal Statistical Society* 109 (1946): 53–7; John Banzhaf, "One Man, 3,312 Votes: A Mathematical Analysis of the Electoral College," *Villanova Law Review* 16 (1968): 304–32. See also Dan Felsenthal and Moshé Machover, *The Measurement of Voting Power: Theory and Practice, Problems and Paradoxes* (Cheltenham: Edward Elgar Publishing, 1998); Annick Laruelle and Federico Valenciano, *Voting and Collective Decision-Making: Bargaining and Power* (Cambridge: Cambridge University Press, 2008).

16. See Banzhaf, "One Man, 3,312 Votes" for the classic treatment of voting power under the Electoral College.

17. One might object that this model is psychologically unrealistic, since voters do not typically vote at random. But the goal of an election model is not to describe how voters in fact reach their voting decisions but, rather, to guide us in estimating the respective probabilities of various different patterns of votes. Psychologically unrealistic models can still generate accurate estimates.

18. In the social choice literature, the assumption of impartial culture is sometimes used to facilitate comparison of different voting rules. See, for example, Vincent Merlin and Dominique Lepelley, "Scoring Run-Off Paradoxes for Variable Electorates," *Economic Theory* 17 (2001): 53–80; or William Gehrlein and Dominique Lepelley, *Voting Paradoxes and Group Coherence: The Condorcet Efficiency of Voting Rules* (Berlin: Springer Science & Business Media, 2010).

19. See Appendix B for proof.

20. This fact is sometimes called Penrose's *square-root law*. See Penrose, "Elementary Statistics of Majority Voting."

winning. And in practice, this is rarely the case. To handle situations where one candidate is favored, a more general *binomial model* can be used. Under a binomial model, an N -voter election is modeled as N tosses of a *biased* coin, where the coin's bias is fixed by the specifics of the case. For example, if Daisy is projected to earn 52 percent of the vote, we can represent each voter's decision as an independent toss of a biased coin which has a 52 percent chance of landing in Daisy's favor.²¹ (Notice that what we were calling random voting, which uses a fair coin, is a special case of this more general binomial model.)

This binomial model has been widely used to estimate d .²² As we saw above, when there is total parity between the candidates (that is, when both candidates are assumed to be exactly equally likely to win), d turns out to be quite high, under the binomial model. But crucially, the value of d falls drastically if we depart from this parity assumption even slightly.²³ If either candidate is assumed to possess a slight advantage over the other, d turns out to be substantially lower, typically far below $1/N$. In Brennan's example discussed earlier, the posited advantage for

21. The motivating idea, here, is that if Daisy is projected to earn 52 percent of the vote, then for an arbitrary citizen, our credence that she will vote for Daisy should be 52%. Accordingly, it seems legitimate to represent each voter's decision as an independent coin toss with a 52 percent chance of landing for Daisy.

22. Authors who employ this model to estimate d include: Banzhaf, "One Man, 3.312 Votes"; Douglas Rae, "Decision Rules and Individual Values in Constitutional Choice," *American Political Science Review* 63 (1969): 40-56; Nathaniel Beck, "A Note on the Probability of a Tied Election," *Public Choice* 23 (1975): 75-80; Guillermo Owen and Bernard Grofman, "To Vote or Not to Vote: The Paradox of Non-voting," *Public Choice* 42 (1984): 311-25; Dennis Mueller, *Public Choice* (Cambridge: Cambridge University Press, 1989); Brennan and Lomasky, *Democracy and Decision*; Alan Carling, "The Paradox of Voting and the Theory of Social Evolution," in *Preferences, Institutions and Rational Choice*, eds. Keith Dowding and Desmond King (Oxford: Clarendon Press, 1995); and Keith Dowding, "Is It Rational to Vote? Five Types of Answer and a Suggestion," *British Journal of Politics and International Relations* 7 (2005): 442-59. This model also is used in Brennan's vivid example, which will be discussed later. For detailed empirical criticism of this binomial model, see Howard Margolis, "The Banzhaf Fallacy," *American Journal of Political Science* 27 (1983): 321-26; Andrew Gelman, Jonathan Katz, and Francis Tuerlinckx, "The Mathematics and Statistics of Voting Power," *Statistical Science* 17 (2002): 420-35; Andrew Gelman, Jonathan Katz, and Joseph Bafumi, "Standard Voting Power Indexes Do Not Work: An Empirical Analysis," *British Journal of Political Science* 34 (2004): 657-74. As we will see in Section V, however, the binomial model is inadequate for a simple reason: it vastly underestimates the probability of an upset.

23. This particular observation is emphasized by Brennan and Lomasky, *Democracy and Decision*, as well as Brennan, *The Ethics of Voting*, 19.

the leading candidate is very slight: 50.5 versus 49.5 percent. But as we saw, the chance of casting a decisive vote in such a case, according to the binomial model, is less than 1 in $10^{2,659}$, which is unfathomably tiny, and obviously far below $1/N$.

So in brief, under an influential and popular election model, d turns out to be many orders of magnitude smaller than $1/N$ under all but the rarest of conditions. This is, I take it, the sort of consideration that leads Brennan and others to conclude that the consequentialist case for voting is hopeless.

But this conclusion is premature. The binomial model is not the only game in town. Some authors have proposed alternative ways to model large elections,²⁴ and others have endeavored to estimate the value of d empirically (i.e., by examining election data and then observing just how large the margins of victory have really been).²⁵ The estimates of d resulting from these alternative approaches all turn out to be considerably greater than the binomial model predicts—indeed, many suggest that on the order of $1/N$.

Given the disagreement found among experts surrounding these issues, it would be convenient if there were a straightforward and uncontroversial way for us to make progress toward estimating d directly, without relying on contentious modeling assumptions, and without analyzing election data in detail. As it happens, progress of this sort can be made. We'll see that, given two weak modeling assumptions—assumptions which any

24. See Gordon Tullock, *Toward a Mathematics of Politics* (Ann Arbor: University of Michigan Press, 1967); William Riker and Peter Ordeshook, "A Theory of the Calculus of Voting," *American Political Science Review* 62: 25–42; I.J. Good and Lawrence Mayer, "Estimating the Efficacy of a Vote," *Behavioral Science* 20 (1975): 25–33. Howard Margolis, "Probability of a Tie Election," *Public Choice* 31 (1977): 134–7; and Gary Chamberlain and Michael Rothchild, "A Note on the Probability of Casting a Decisive Vote," *Journal of Economic Theory* 25 (1981): 152–62.

25. Estimating d empirically is a delicate task, since few large elections, if any, have actually been decided by a single vote. Accordingly, it might seem that our sample size is too small to support a trustworthy estimate. Some authors explore an interesting way around this problem. They tally up the number of elections where the margin of victory was very small—say within 100 votes. Each such election counts, for estimation purposes, as $1/100$ th of a case where a decisive vote was cast. See Andrew Gelman, Gary King, and W. John Boscardin, "Estimating the Probability of Events that Have Never Occurred: When Is Your Vote Decisive?" *Journal of the American Statistical Association* 93 (1998): 1–9; Gelman, Katz, and Tuerlinckx, "The Mathematics and Statistics of Voting Power"; and Casey Mulligan and Charles Hunter, "The Empirical Frequency of a Pivotal Vote," *Public Choice* 116 (2003): 31–54, for instances of this strategy.

plausible model should make—the probability of casting a decisive vote, d , can be shown to be greater than $1/N$ as long as both candidates have at least a 10 percent chance of winning.

B. Warm-Up Example

Before we look at the argument in detail, it will be helpful to examine a warm-up example first, to gain an intuitive sense of how the argument will work.

For ease of presentation, let's suppose that the electorate is composed of you and one million others—so $N = 1,000,001$. And to make things as tough on ourselves as possible, let's suppose that all of your fellow citizens are certain to vote—no one will abstain, except maybe you. Finally, let's suppose that Daisy and Donald, our two candidates, are exactly equally likely to win. (Later we will relax this assumption.) Given this background, how likely is it that your vote will be decisive? Likelier than $1/N$, I think. This can be shown by way of a simple counting argument. Here are the different ways that *everyone else* could have voted.

	Votes for Donald	Votes for Daisy	Outcome
1,000,001 possibilities in total	1,000,000	0	Donald wins regardless of what you do
	999,999	1	Donald wins regardless of what you do
	999,998	2	Donald wins regardless of what you do

	500,001	499,999	Donald wins regardless of what you do
	500,000	500,000	Your vote is decisive
	499,999	500,001	Daisy wins regardless of what you do

	2	999,998	Daisy wins regardless of what you do
	1	999,999	Daisy wins regardless of what you do
	0	1,000,000	Daisy wins regardless of what you do

Notice that your vote is guaranteed to be decisive in exactly one case: the case where everyone else ends up tied. Since there are 1,000,001 outcomes in total, the probability that your vote is decisive will depend upon the relative likelihood of the middle row.

If, for some reason, each of these outcomes were equally likely, then the probability that your vote is decisive would be exactly $1/1,000,001$ (which is precisely $1/N$). But presumably, these outcomes are not all equally likely, for we know that the two candidates are equally likely to win. In light of this, the outcomes toward the top and bottom of our list are less likely, and the outcomes toward the center are, presumably,

likelier. If this is right, the chance that your vote is decisive seems greater than $1/1,000,001$ —and probably much greater.

Now admittedly, the foregoing reasoning depends crucially on the assumption that our candidates are exactly equally likely to win—an assumption which is rarely true. But as we'll see, a version of the argument can be advanced even if this assumption is relaxed.

C. Two Modeling Assumptions

To present the more general argument for the chances condition, we will need to make two important modeling assumptions. These assumptions pertain to the proportion of votes earned by the leading candidate, *when we count up everyone's votes except for your own*.

Partial Unimodality: The leading candidate is at least as likely to earn exactly half of the vote as she is to earn any precise share of the vote smaller than this.²⁶

Narrow Upsets: If the leading candidate fails to earn a majority, then the likelihood that she comes within ten percentage points of her opponent is at least $\frac{1}{2}$.

Both assumptions are true of every election model with which I am familiar, including the binomial model favored by Brennan. Nonetheless, I will explain and defend each assumption in turn.

Partial Unimodality: Partial Unimodality asserts that, when we tally up everyone's votes except for your own, the leading candidate is at least as likely to earn exactly 50 percent of the vote as she is to earn any particular smaller share of the vote. Why think that this is true?

For starters, here's a quick and intuitive defense of the idea. Suppose that Daisy is ahead of Donald in the polls. Presumably, there is some share of the vote greater than 50 percent, which Daisy is projected to receive—this might be her polling average. No matter what this projected outcome is, we know that the outcome in which Daisy earns 50 percent of the vote is *closer* to this projected outcome than are all outcomes in which Daisy earns less than 50 percent. Other things equal, it is reasonable to

26. This formulation of Partial Unimodality assumes that N is odd. The even case is handled in Appendix A.

suppose that outcomes which are closer to the projected outcome are at least as likely as those which are farther away.

This intuition can be spelled out more thoroughly, though. Partial Unimodality follows from a more basic modeling assumption, often called *unimodality*. When a probabilistic model is unimodal, it means that the likeliest outcomes are clustered together, while outcomes become less and less likely as one moves further and further from that cluster.

This more general assumption of unimodality, as a constraint on election models, can be defended on Bayesian grounds. Suppose we're trying to estimate the degree of support for Daisy among the general public. We conduct a random sample and find that, among those polled, 53 percent support Daisy. Under what conditions would this finding be likeliest? Well, the finding would of course be likeliest if the true degree of support were 53 percent; the finding would be slightly less likely (but still unsurprising) if the true degree of support were 52 or 54 percent; and the finding would be very improbable indeed if the true degree of support for Daisy were, say, 30 or 80 percent. In general, our finding of 53 percent was to be expected *to the extent that 53 percent is close to the true degree of support for Daisy among the public*. This observation, on a standard Bayesian picture, corroborates a unimodal assignment of probabilities to outcomes centered at 53 percent (or, visually, a hill-shaped distribution with a "peak" at 53 percent).²⁷

Narrow Upsets: Narrow Upsets asserts that *if* Daisy fails to earn a majority, then the chance that she comes within ten points of her opponent is at least $\frac{1}{2}$. Why think that this is true?

Leading candidates do lose from time to time. Election forecasts are not perfectly accurate. But when leading candidates do lose, they tend to do so narrowly: Election forecasts are rarely wildly off the mark. In case of an *upset* (an unexpected victory), it is common for the losing candidate to lose by a narrow margin, maybe a percentage point or two. Given this fact, the proposed modeling assumption says something very cautious: only

27. Admittedly, one can cook up cases where unimodal expectations would not be appropriate: Suppose that two polls are conducted, and you suspect that exactly one is fraudulent, but you do not know which. If the two polls estimate the degree of support for Daisy at 55 and 45 percent, respectively, then you should expect that the true level of support for her is either at ~ 55 or ~ 45 percent, but probably not 50 percent. However, if both polls were honestly and competently conducted (and sampled equally many citizens), then the likeliest explanation is simply that the true level of support for Daisy is at ~ 50 percent.

that, if the leading candidate fails to earn a majority, then there's at least a $\frac{1}{2}$ chance that the leading candidate comes within *ten* percentage points of her opponent.

One can consult historical data for extra assurance here. Looking back on U.S. Senate races from the last ten years, one finds that in cases where the leading candidate lost, he/she still came within ten points of winning about 92 percent of the time.²⁸ This is considerably greater than the proposed bound of $\frac{1}{2}$.

D. The Argument

With these modeling assumptions in place, it can be argued that the value of d is often going to exceed $1/N$, so long as the election is reasonably competitive.

For illustration, let's again suppose that the electorate is composed of you and one million fellow citizens, all of whom are certain to vote ($N = 1,000,001$).²⁹ But this time, we will not assume that Daisy and Donald are equally likely to win. Instead, we'll assume that while Daisy is the leading candidate, Donald still has *at least a 10 percent chance of achieving an upset victory*. From these limited assumptions, what can be inferred about your chance of casting a decisive vote?

Well, we know that there's clearly at least a 10 percent chance that Daisy fails to earn a majority. Given Narrow Upsets, we know that *if* Daisy fails to earn a majority, then there's at least a $\frac{1}{2}$ chance that she comes within ten percentage points of Donald. Putting these claims together, we obtain the observation that there's at least a 5 percent chance that Daisy fails to earn a majority but still comes within ten percentage points of Donald—or in other words, there's at least a 5 percent chance that *Daisy earns between 45 and 50 percent of the vote*. This observation, together with Partial Unimodality, already entails that $d \geq 1/N$. To see this, consider the

28. Senate races from 2010, 2012, 2014, 2016, and 2018 were analyzed. The RealClearPolitics' (<https://www.realclearpolitics.com/>) polling average was used to identify leading candidates. Upsets occurred in Colorado 2010, Nevada 2010, Montana 2012, North Dakota 2012, Kansas 2014, North Carolina 2014, New Hampshire 2016, Pennsylvania 2016, Wisconsin 2016, Arizona 2018, Florida 2018, and Indiana 2018. Of these, only in the Kansas 2014 race was the margin of victory greater than ten percentage points.

29. See Appendix A for a more general version of the argument, which does not assume N to be any particular number.

following list, which describes the different ways that Daisy could receive between 45 and 50 percent of the vote (not counting you).

Votes for Donald	Votes for Daisy	Daisy's share	Outcome
549,999	450,001	45.0001%	Donald wins regardless of what you do
549,998	450,002	45.0002%	Donald wins regardless of what you do
549,997	450,003	45.0003%	Donald wins regardless of what you do
...
500,002	499,998	49.9998%	Donald wins regardless of what you do
500,001	499,999	49.9999%	Donald wins regardless of what you do
500,000	500,000	50.0000%	Your vote is decisive

50,000 possibilities in total

We want to know how likely the outcome at the bottom is. That's the one case where your vote will be decisive. We have already inferred that there is at least a 5 percent chance that the actual outcome appears somewhere in this list. But how likely is the outcome at the bottom, in comparison to the others?

If, somehow, the outcomes in this list were all equally likely, then the probability of your casting a decisive vote would be at least $(5 \text{ percent}) \times (1/50,000) = 1/1,000,000$, which is already greater than $1/N$. But given the assumption of Partial Unimodality, things look even better, for we know that the outcome at the bottom of the list is *at least as likely* as everything above it. So the probability of your casting a decisive vote is $(\text{at least } 5 \text{ percent}) \times (\text{at least } 1/50,000)$, which is at least $1/1,000,000$, which is greater than $1/N$. Even given our very limited assumptions, the chances condition still turns out to be true.

The reasoning outlined above can be used to derive a more general lower bound on d in terms of u , the probability of an upset. Specifically, it can be proven from our modeling assumptions that, for large N , $d \geq 10u/N$.³⁰ This observation helps us assess how d varies, under different conditions of electoral competitiveness.

Probability of upset (percent)	d is at least . . .
50	$5/N$
40	$4/N$
30	$3/N$
20	$2/N$
10	$1/N$

30. See Appendix A for proof.

Unsurprisingly, when an election is more competitive, a greater minimum value for d can be established—approximately $5/N$ for the most competitive elections. But even for quite one-sided elections where one candidate has a 90 percent chance of winning, d can still be shown to be at least $1/N$.

The bottom line: Given two attractive modeling assumptions, the chances condition is true in a wide range of real-world electoral circumstances.

E. Worries About Expected Value Reasoning

The argument is thus complete. We have seen that voting is rational so long as two key conditions are met, and we have examined some grounds for thinking that, indeed, both conditions are often satisfied in the real world. At this point, I would like to address a critical reaction some readers may have, which can be expressed as follows.

At the end of the day, you're telling me it's rational to vote on the hope that my individual vote tips the election. You don't deny that this prospect is incredibly tiny (maybe $1/100,000,000$ for a national election in the United States). Now, I understand the decision-theoretic argument you've given (the stakes are high, the huge stakes get weighed against the tiny chance, etc.), and perhaps the math works out as you say. Even so, I find it difficult to accept that a rational actor would vote on the basis of such a vanishingly small probability. Perhaps it's true that orthodox decision theory construes rationality this way. But must I?

I can appreciate the objector's uneasiness here. When I vote in large elections, I sometimes feel an intuitive sense of pointlessness, and the thought that my act has positive expected value does not entirely quell these doubts. "One in a hundred million? That is just not going to happen, period," I might think to myself. And it's worth noting that, even apart from the voting context, some have raised independent reasons to worry about how decision theory handles tiny probabilities.³¹ There are two points worth making in response, however.

31. See Nick Bostrom, "Pascal's Mugging," *Analysis* 69 (2009): 443–5, for an amusing, if somewhat troubling, example.

First, even if it is conceded that sufficiently tiny probabilities (such as 1/100,000,00) should simply be ignored in all cases, the argument presented in this article still shows that, contrary to popular opinion, voting to change the outcome often has positive expected value. This is a fact worth acknowledging.

Second, while acting on the basis of tiny probabilities may seem troubling, it's not clear that a viable alternative is available. The most natural alternative view in the vicinity simply says that sufficiently tiny chances can rationally be *ignored*, regardless of the stakes. But this view comes with its own difficulties. Parfit discusses one troubling consequence of the "ignore-tiny-chances" view:

It may be objected that it is irrational to consider very tiny chances. [. . .] Suppose that nuclear engineers did ignore all chances at or below the threshold of one-in-a-million. It might then be the case that, for each of the many components in a nuclear reactor, there is a one-in-a-million chance that, in any day, this component would fail in a way that would cause a catastrophe. It would be clearly wrong for those who design reactors to ignore such tiny chances. If there are many reactors, each with many such components, it would not take many days before the one-in-a-million risk had been run a million times. There would fairly soon be a catastrophe. When the stakes are very high, no chance, however small, should be ignored.³²

Parfit's argument, of course, should not necessarily be regarded as the final word on the matter. But it does suggest, at least, that we should not be too quick to abandon the standard picture. And given the standard picture, the conclusion that voting is often rational will be difficult to avoid.

V. WHAT'S WRONG WITH BRENNAN'S EXAMPLE?

At the outset of this article, we discussed an example, which purportedly illustrates that voting to change the outcome is not rational, given standard decision-theoretic assumptions. If it is assumed that there are 122,293,322 voters, that each individual voter independently has a 50.5 percent chance of voting for Daisy, and that electing Daisy over Donald

32. Parfit, *Reasons and Persons*, 74.

according to the binomial model, the chance that Donald wins or ties is less than 8.84×10^{-2653} . This is roughly the chance of a one-in-a-million event happening 442 times in a row.³⁵

The eye-popping figure Brennan's example generates is not proof of the futility of voting to change the outcome; it is a consequence of the example's use of the binomial model, which is a very poor way to model real-world elections. Which model should we use instead? That's a complex and substantial question, which will require detailed work by economists, political scientists, and statisticians to settle. But as we've seen, even without relying on any controversial modeling assumptions, it can be argued that the probability of casting the deciding vote is very often greater than $1/N$, sometimes considerably so. If we return to Brennan's example with this observation in mind, the expected contribution to the public good of a vote for Daisy is nothing like a millionth of a penny—it's at least +\$134.92, and probably much greater.³⁶ All things considered, then, voting seems well worth the cost.

APPENDIX A

ESTIMATING THE CHANCE OF CASTING A DECISIVE VOTE

In Section IV.D, we saw that when $N = 1,000,001$, the conclusion that $d \geq 1/N$ can be derived from our modeling assumptions, so long as the underdog candidate has at least a 10 percent chance of winning. Below, I'll offer a more general version of the argument, discussed briefly at the end of Section IV.D, which establishes a lower bound on d in terms of u , the probability of an electoral upset. Specifically, we'll see that, given generalized versions of our two modeling assumptions, $d \geq 10u/N$, which implies that $d \geq 1/N$ in a wide class of cases.

Generalized Argument

The usual assumptions are in effect: There are N citizens, all of whom will vote for one of the two candidates.

Let N^- stand for the number of citizens, *not including you*.

Let L stand for the number of votes, from that group of N^- citizens, earned by the leading candidate.

35. See Appendix B for proof.

36. $EB_{\text{vote}} = (\frac{1}{2})(B)(d) = (\frac{1}{2})(+\$33 \text{ bil.})(1/N) = (\frac{1}{2})(+\$33 \text{ bil.})(1/122,293,322) \approx +\134.92 .

Let t stand for what might be called a *triggering number*—that is, the number of votes the leading candidate would have to earn, from your fellow citizens, in order for *your* vote to be decisive. (When you're voting *for* the leading candidate, $t = \lfloor N^-/2 \rfloor$; when you're voting *against* the leading candidate, $t = \lceil N^-/2 \rceil$.)

Let S stand for the set: $\{s \in \mathbb{N} \mid t - \lfloor N^-/20 \rfloor < s \leq t\}$. (When $L \in S$, the leading candidate has earned roughly between 45 and 50 percent of the vote.)

Let u stand for the *ex ante* probability that the leading candidate earns the triggering amount or fewer votes from your fellow citizens. The more competitive the election, the greater u will be.

Here are generalized versions of our two modeling assumptions.

Partial Unimodality (generalized): $\forall n < t [P(L = t) \geq P(L = n)]$.

Narrow Upsets (generalized): $P(L \in S \mid L \leq t) \geq 1/2$.

When N is large, these assumptions assert almost exactly what the original ones did. (When N is small, they do not; for example, when $N \leq 20$, generalized Narrow Upsets is automatically false, for S is empty.) They are formulated in a more general manner so that they can be true whether N is even or odd, and no matter which candidate you are voting for.

Given these assumptions, the claim that $d \geq 10u/N$ can be proven as follows:

$$\begin{aligned}
 d &= P(L=t) \\
 &= P(L=t \mid L \in S) \times P(L \in S) \\
 &= P(L=t \mid L \in S) \times P(L \in S \mid L \leq t) \times P(L \leq t) \\
 &\geq_i P(L=t \mid L \in S) \times 1/2 \times P(L \leq t) \\
 &\geq_{ii} (20/N^-) \times 1/2 \times u \\
 &= 10u/N^- \\
 &> 10u/N
 \end{aligned}$$

(inequality *i* follows from generalized Narrow Upsets; inequality *ii* follows from generalized Partial Unimodality, together with the observation that $|S| \leq N^-/20$.)

As we saw in the text, this lower bound on d implies that $d \geq 1/N$ under a wide variety of electoral circumstances.

APPENDIX B

TROUBLE FOR THE BINOMIAL MODEL³⁷

In Section V, we discussed difficulties associated with the binomial model, which models an N -voter election as, in effect, a series of N independent coin tosses. The *bias* of the coin can be varied to fit the specifics of the case. So if both candidates are equally likely to win, the imagined coin would be a fair one, equally likely to produce a vote for either candidate on any given toss. But if a given candidate is favored, a *biased* coin is used instead. We'll let r stand for the coin's bias toward the favored candidate (that is, for the probability that the imagined coin produces a vote for the favorite on a given toss).

According to the binomial model, the probability that the leading candidate receives exactly X votes, $P(L_X)$, is given by the following expression:

$$\Pr(L_X) = \binom{N}{X} \cdot r^X \cdot (1-r)^{N-X}$$

The term $\binom{N}{X}$ typically read as “ N choose X ” stands for the number of different sets of size X which can be chosen from a set of size N . It is defined as follows:

$$\binom{N}{X} = \frac{N!}{X!(N-X)!}$$

In the present context, the “ N choose X ” term represents the number of different combinations of X voters (chosen from our N -voter population) that could conceivably team up to vote for the favored candidate.

For illustration, let's consider an election with 500,000 voters ($N = 500,000$). The two candidates are equally likely to win ($r = 0.5$). On the binomial model, how likely it is that the election ends in an exact tie? In other words, what's $\Pr(L_{250,000})$?

37. Thanks to an Associate Editor at *Philosophy & Public Affairs* whose suggestions improved the readability of this section.

$$\Pr(L_{250,000}) = \binom{500,000}{250,000} \left(\frac{1}{2}\right)^{250,000} \left(1 - \frac{1}{2}\right)^{250,000} \approx 0.0011284.$$

The likelihood of such an outcome turns out to be about 1 in 1,250. The calculation above (as well as those that follow) cannot realistically be carried out by hand. Computational software can be used to approximate the value of the desired expression(s). For this example, the reader can obtain the result stated above with the following input query using WolframAlpha (<https://www.wolframalpha.com/>):

$$(500,000 \text{ choose } 250,000)((1/2)^{25,0000})((1 - 1/2)^{250,000}).$$

So in a nutshell, that's how the binomial model can be used to estimate the probabilities of certain electoral outcomes.

In the text, we saw that this model makes some peculiar predictions. Before we return to those, I want to consider a smaller-scale example, which will illustrate the inadequacy of the binomial model more easily, and in a way that will be easy for readers to verify for themselves.³⁸

Smaller-Scale Example: Suppose that there are 500,000 voters ($N = 500,000$) and, as is in Brennan's example, the chance of a given voter's voting for Daisy is 50.5 percent ($r = 0.505$). Using the binomial model, let's estimate the probability that the underdog candidate, Donald, wins or ties.

$$\begin{aligned} \Pr(\text{Donald wins or ties}) &= \Pr(L_0 \vee L_1 \vee \dots \vee L_{250,000}) \\ &= \Pr(L_0) + \Pr(L_1) + \dots + \Pr(L_{250,000}) \\ &= \sum_{i=0}^{250,000} \binom{500,000}{i} (0.505^i) (0.495)^{500,000-i} \\ &\approx 0.000000000000775235. \end{aligned}$$

38. The sheer magnitude of the electorate in Brennan's example ($N = 122,293,222$) makes calculation significantly more cumbersome, even with the aid of a computer. To justify the claims made in the text, we will need to use algebraic tricks and properties of inequalities to make the expressions more manageable—and even then, specialized software is needed for the computational part. But I thought it advisable to include at least one example which could be easily verified (e.g., using WolframAlpha) so that the inadequacy of the binomial model would be apparent to all readers.

