

Utilitarianism and Voting

Zach Barnett

Guest essays represent only the views of the author(s).

Table of Contents

1. Introduction
2. The Chance of Casting a Decisive Vote
3. The Argument (applied to our case)
4. Impartiality and Altruism
5. Determining the Better Candidate
6. Further Reading

The essay below, which aims to quantify the social benefits of voting in utilitarian terms, is intended as a complement to existing resources, such as [“If you care about social impact, why is voting important?”](#) from *80,000 Hours*.

Introduction

This essay discusses the value of voting from a utilitarian perspective. To some, this may seem like the wrong place to start. After all, some of the most popular rationales for voting are not at all utilitarian. For example, some say you should vote to...

- fulfill your civic duty,

- express your preference,
- share in a winning effort,
- avoid being complicit in wrongs your government may commit.

These popular rationales have something in common. They can apply even *if your vote has no chance of affecting the result*. From a certain angle, this can seem unsatisfying.

Imagine you are one of three students eligible to vote for an important policy change in your political philosophy class. Before you go vote, you happen to see your classmates, the two other eligible voters. With a smile, they inform you that they've already both voted—and that they voted for the same thing.

Welp. Now voting looks somewhat pointless. Your vote has no chance of affecting the result. But imagine that your classmates try to encourage you:

Hey! You should totally still go vote—it's your civic duty after all! Don't you want to express your preference? By voting with us, you could share in a winning effort! And by voting against us, you'll avoid being complicit in any wrongs we may have committed.

I think we'd have to forgive someone who was unmoved by these arguments.

The utilitarian approach is different—and, to my mind, potentially more satisfying. Should you vote? Utilitarianism says: *it depends*. Would voting, rather than staying at home, tend to promote well-being? If so, vote! If not, don't bother. (Of course, this is how utilitarianism answers *any* question about whether you should do something.)

But—some might worry—does voting tend to promote well-being in real life? It's tempting to think that it doesn't promote well-being very much. For your vote can only affect the result of the election if it is decided by a single vote, and, for large elections, the chance that it literally comes down to a single vote is tiny.

Concretely, Jason Brennan estimates that the expected benefit of your vote, in dollars, is only $\$4.77 \times 10^{-2650}$, even if the election is assumed to be competitive and even if your preferred candidate is assumed to be worth \$33 billion more to the common good than their opponent.¹

But Brennan's figure is based on an estimate of the probability of casting a decisive vote that is off by thousands of orders of magnitude. (For one, the model from which this defective figure is derived implies that the probability of an electoral upset is 8.84×10^{2653} , which is as likely as a one-in-a-million event happening 442 times in a row.²) When we replace this inaccurate estimate with a better one, we obtain a figure that renders voting worthwhile in a wide range of cases.

In what follows, I walk through how to obtain this better estimate, while making very few assumptions about the world. According to this better estimate, the chance that your vote is decisive is often greater than $1/N$, where N is the number of voters. When this estimate holds, the expected value of voting is on a par with the *average* benefit of installing the better candidate. So however good you think it would be, per person, to have the better candidate in charge—voting for that candidate, in a competitive election, has an expected value on a par with helping one person that much.

After, we'll discuss two reasons to worry about this utilitarian defense of voting—one deriving from worries about how often voters possess the altruistic values needed for the argument to work, and another deriving from worries about your reliability in identifying the better candidate.

The Chance of Casting a Decisive Vote

There are two candidates on the ballot: the Favorite and the Underdog. The election is decided by majority rule—so the candidate who gets the most votes wins. And there are exactly 1,001 voters including you. With these simplifying

assumptions in place, your vote will be decisive if *everyone except for you* ends up in a 500 to 500 tie. But what's the chance of that?

When we tally up everyone else's votes, there are various possible scores. Consider a graph that represents the likelihood of all of these different Favorite-to-Underdog scores. If the graph were completely flat, it would mean that all 1,001 possible scores are exactly equally likely. If the graph were an upward-sloping line, it would mean that the probability of an outcome is directly correlated with the number of votes the Favorite receives, with a 1000-0 victory being the likeliest.

Obviously, these examples are very unrealistic. But what should the graph look like? There are many models out there, but we don't want to commit ourselves to anything too specific. What we can assume, though, is that the shape of the curve will usually be approximately hill-shaped—centered around some projected outcome. When a graph has this property—when it goes up and then down—it is said to be *unimodal*.

Virtually all existing models have this property. For example, see [this](#) primer from Nate Silver, explaining the basics of [FiveThirtyEight](#)'s election forecasts. The models discussed there are all unimodal.

But let's not rest on the authority of others. Say you're trying to estimate how many red and blue M&Ms there are in a massive collection of millions of M&Ms. You do a random sample, and you find that, in your sample, the ratio of red to blue is 55/45. That finding is likeliest if the true ratio is 55/45. That finding is a bit less likely if the true ratio is 56/44, and less likely still if the true ratio is 60/40. Broadly speaking, something similar is true for political polling. Say that the polls are showing a competitive race—maybe 51 to 49. Other things equal, a reasonable projection for the result is that it will roughly match the polls. Other outcomes are less likely to the extent that they deviate from this projection. So our first assumption is that **(1) the graph is unimodal.**

The second assumption we'll need is that the election is reasonably competitive—or specifically, that **(2) *the trailing candidate, Underdog, has at least a 15% chance of winning.***

This assumption won't always hold. If the election is very one-sided, voting won't be especially valuable. So let's restrict our attention to competitive elections.

(Someone might wonder, here, about uncompetitive *states* in an electoral college system. Under an electoral college system, voters in uncompetitive states have a far smaller expected impact than those in competitive states. This is a bad feature of the electoral college system, and is one reason to prefer moving to a different one, which distributes voting power more equitably.)

The third and final assumption we'll need is that the Favorite is especially unlikely to get blown out—or specifically, that **(3) *the probability that the Underdog wins by 20+ percentage points is less than 1%.***

This assumption is extremely cautious. In a large election with lots of data available, it's almost unfathomable that a candidate would be projected to win and then end up with a vote share *in the 30s or below*. Errors of this magnitude almost never happen. But we're not even assuming that this won't happen—we're just assuming that it has less than a 1% chance.

So we have three pretty weak assumptions: (1) a unimodal graph; (2) a competitive election; (3) and a less than 1% chance that the Underdog wins by 20+ percentage points. From these weak assumptions, it can be *proven* that the chance of a 50-to-50 tie is greater than $1/N$, which is $1/1001$ in this case.

The Argument (applied to our case)

Distinguish three possible outcomes:

[A] The Underdog receives less than 500 votes (and loses).

[B] The Underdog receives between 500 and 600 votes (and wins or ties).

[C] The Underdog receives more than 600 votes (and wins by 20%+).

Since it's a competitive election (assumption 2), [A] has at most an 85% chance. That leaves 15%+ for [B] and [C] to share. And [C] is the unlikely scenario where the polls are massively wrong, which has at most a 1% chance (assumption 3).

That leaves a 14%+ chance of [B]—the Underdog receives between 500 and 600 votes. 101 distinct possibilities belong to this second outcome. If they all were equally likely, then each would have a probability of $.14 \times 1/101 = 7/5050 \approx 1.4/1001 = 1.4/N > 1/N$. But by unimodality (assumption 1), the graph is upward-sloping in this interval. So the 500-to-500 tie possibility is the likeliest of these outcomes. Thus, its probability is at least $7/5050$, and probably much greater. So we're done.

Importantly, nothing about the argument was sensitive to the specific number of voters that was chosen. It holds equal if there are a million voters, or a hundred million. As long as these minimal conditions hold, your chance of casting a decisive vote is actually higher than your chance of being chosen at random in a lottery from among the electorate!

Given the high stakes associated with many elections, **this probability is plenty high enough to make voting relatively impactful in expected value terms.** For example, using Brennan's assumptions together with the foregoing estimate, we obtain the result that your vote has an expected contribution to the common good of at least \$187.20.³

More generally, if breaking a tie in favor of your preferred candidate is worth B to society as a whole, then it follows that your vote is worth at least B/N , where N is the number of voters. It's helpful to think of this expression, B/N , as the average, per-person benefit of having your preferred candidate in charge. Whatever that average benefit is, we can conclude that voting in a competitive election is, in the long run, as valuable as helping one person exactly that much.

Impartiality and Altruism

So far, we have assumed that you, the voter, are utilitarian in your values. If we relax that assumption—if we assume that you care more about yourself—then the argument loses some of its force. The worry is nicely summarized by

[Michael Huemer](#):

Some argue that if you are a pure altruist (you value a random stranger's interests equally with your own), then the expected benefit of voting can still be larger than the cost (in terms of the time & effort it takes to vote). This is correct. The probability of affecting the outcome is very small, but if it happens, you'll have a large total effect on society, since the President affects a lot. However... This only works if you are way more altruistic than, I believe, any actual human being is. I think the typical human values his own life at least thousands of times more than that of a stranger, and maybe millions or billions of times more.

Of course, this worry does not necessarily undermine the *moral* case for voting. But it threatens to undermine the *rational* case for voting, since it suggests that no one will actually have the relatively impartial motivational profile the argument assumes.

Since you're currently reading an essay on utilitarianism.net, you may be somewhat more optimistic about the possibility of pure altruism than Huemer is. But even if we grant that no human being is perfectly altruistic all of the time, I believe that people often strive to (and sometimes do) engage in politics in broadly altruistic ways, which are reflective of giving equal consideration to everyone's interests.

For example, consider a policy you support. Presumably, you support that policy not because it would benefit you specifically, *but because you think it would be good for society as a whole*, in some sense. I can't deny that I'm selfishly motivated some of the time, but it seems implausibly cynical to deny

that I can have sincere political convictions based on principles other than self-interest.

All this suggests that—at least for some of us, some of the time—we have the capacity to engage politically in a way that is impartial and reflective of giving everyone’s interests equal consideration, even if we don’t always adhere to those values. In short, we can be temporarily altruistic, while thinking through our political views and going to the voting booth. If this is right, then Huemer’s dismissal is premature.

Before moving on, it’s worth briefly emphasizing the specific values that flow from a utilitarian ideology. For while some voters might favor policies that are good for their neighborhood, for their state or territory, for their own demographic group, or even for their country as a whole, utilitarians are committed to promoting the well-being of any creature that can suffer. So utilitarian voters will take into account the interests of the global poor, non-human animals, and future generations, in their political decisions.⁴

Determining the Better Candidate

Even if altruism is assumed, though, there is still one other worry about voting worth taking seriously: [What if you’re voting for the worse candidate?](#)

So far, we have uncritically assumed that, of course, you’ll be voting for the good guy or gal. But this isn’t necessarily true. There are well-meaning voters on both sides of the political aisle, who truly believe that their preferred candidate will promote well-being better than the alternative. In fact, there’s a lot of research showing just how ill-informed partisan voters typically are, including those who are well-meaning and well-educated.

In one experiment, they asked people whether they thought taxes on those earning \$250,000 or more should go up, go down, or stay the same.

Overwhelmingly, they said up. Definitely Up. Then they asked what tax rate they thought it should go up to. And most of those people chose a tax rate that

was below the prevailing rate at the time!⁵ It's hard to have a reasonable view on whether something should go up or down if you have no idea what the current value is.

A second finding, due to Diana Mutz, is even more telling. Mutz found that the better someone is at explaining the views of the opposing side, the *less* likely they are to be politically active. Conversely, those who participate in politics the most, on average, tend to be *least able* to articulate the opposing side's motivations.⁶

There are many other findings like these.⁷ What should we make of them? It's tempting to dismiss them. It's tempting to think that, surely, these sorts of biases afflict others, *but not me*. Such a reaction, though, is defensible only if there is evidence that you are unusually informed, perceptive, and unbiased. And most of us do not possess this kind of evidence.

Optimistically, though, I'm inclined to think that it is sometimes possible to be aware of and—to some extent—compensate for known distorting biases, such as self-interest and tribalism. And I believe that there are at least some followers of politics who are relatively reliable in their judgments. But people who fit this description may be rarer than we care to admit.

The bottom line? There are reasons to worry about the reliability of our own political judgments. But this needn't lead to disengagement; rather, it should motivate us to exercise greater care in thinking about politics. And when we do the math, what we discover is that, under favorable conditions, voting has a meaningful impact—one that is positive and significant.

About the Author

[Zach Barnett](#) is an Assistant Professor of philosophy at the University of Notre Dame, where he teaches courses in ethics, epistemology, and decision theory.



How to Cite This Page

Barnett, Z. (2024). Utilitarianism and Voting. In R.Y. Chappell, D. Meissner, and W. MacAskill (eds.), *An Introduction to Utilitarianism*, <<https://www.utilitarianism.net/guest-essays/utilitarianism-and-voting>>, accessed 9/18/2024.

Further Reading

- *Hearing the Other Side: Deliberative vs. Participatory Democracy* by Diana Mutz (2006)
- *The Myth of the Rational Voter: Why Democracies Choose Bad Policies* by Bryan Caplan (2007)
- [Voting as a Rational Choice: Why and How People Vote to Improve the Well-being of Others](#) by Aaron Edlin, Andrew Gelman, and Noah Kaplan (2007)
- *The Ethics of Voting* by Jason Brennan (2011)
- *Free to Move: Foot Voting, Migration, and Political Freedom* by Ilya Somin (2013)
- [Why You Should Vote to Change the Outcome](#) by Zach Barnett (2020)
- *Why It's OK to Ignore Politics* by Christopher Freiman (2020)
- [If you care about social impact, why is voting important?](#) by Robert Wiblin (2020)

1. This estimate is drawn from *The Ethics of Voting* by Jason Brennan (pp. 18–20). ↩
2. To see in greater depth how the specific ‘binomial model’ employed by Brennan leads to strange and implausible consequences, see “[Why You Should Vote to Change the Outcome](#),” by Zach Barnett. ↩
3. This figure is derived from Brennan’s assumptions—that the candidate you’re voting for is 33 billion dollars better for society than her rival and that there are 122,293,322 voters (as in the 2004 US presidential election). The argument given above entails that your chance of being decisive—provided that the assumptions are all satisfied—is at least $\sim 1.4/N$. Since breaking a tie in favor of your preferred candidate is worth 16.5 billion dollars, the expected value of voting is estimated at the product of these values, which leads to the figure shown in the text. ↩
4. I’m grateful to Darius Meissner for suggesting that I emphasize these points. ↩
5. This finding is discussed in Ilya Somin’s (2007) book *Democracy and Political Ignorance* (pp. 17–18). ↩
6. Diana Mutz’s (2006) book *Hearing the Other Side* discusses these sorts of effects. ↩
7. Bryan Caplan’s (2007) book *The Myth of the Rational Voter* surveys a lot of these results. ↩