

Expert Disagreement and the Duty to Vote
Devin Lane
University of North Carolina, Chapel Hill
(forthcoming in *Philosophers' Imprint*; please cite final version)

In this paper, I argue that widespread expert disagreement about sufficiently many issues central to a democratic decision-making procedure can nullify the duty to vote. I begin by drawing a distinction between different ways that we might conceive of the duty to vote, i.e. whether it is a duty to vote, no matter how one votes, or a duty to vote well. I then review some prominent arguments in favor of the existence of the duty to vote and suggest that they go through only when the potential voter in question would vote well. So the most promising existing arguments for the duty to vote, if successful, really show not merely that citizens have a duty to vote, but more strongly that they have a duty to vote well. I then discuss the epistemology of expert disagreement to show how circumstances of expert disagreement can nullify such a duty by making it impossible to vote well. To my knowledge, no one has yet brought considerations of the epistemology of expert disagreement to bear on this topic. That is my aim in the main argument of this paper. The remainder of the paper addresses alternative characterizations of ‘voting well’, and argues that on other plausible characterizations, expert disagreement will still often nullify the duty to vote.

1. Introduction

There has recently been much debate about whether citizens in democracies have a duty to vote. Those who think there is such a duty include Julia Maskivker (2019), Zach Barnett (2020), Alex Guerrero (2015), and Eric Beerbohm (2012). Those who think there is no such duty include Jason Brennan (2011) and Loren Lomasky and Geoffrey Brennan (2010).

There are three related but distinct (potential) duties that are at issue in this debate, which are not always well-distinguished. One is the (potential) duty to vote simpliciter. One satisfies this duty whenever one votes, regardless of how one votes. The second is a (potential) *conditional* duty that *if*

one votes, then one must vote well. This duty is violated only by someone who *does* vote, but does not vote well: for example, by voting randomly, or on the basis of unjustified beliefs.¹ The third is a (potential) outright duty to vote well. This is violated *both* by someone who does not vote, and by someone who votes but does not do so well. Very plausibly, this third duty follows from the first two. If we ought to vote, and we are required to vote well if we vote, then we ought to vote well.

I will argue that there is not an outright duty to vote well, at least under certain realistic circumstances, because the duty would be impossible to discharge. From this it follows that either there must be either be no conditional duty to vote well or no outright duty to vote. I will argue that it is the outright duty to vote that should (in these cases) be rejected.

The circumstances under which I think that the duty to vote is nullified are conditions of *expert disagreement*, though I need to clarify what this means. Consider an election that centers on ten issues. Suppose that there is expert consensus about nine of them, all reasonably taken to be highly important by the citizen considering whether to vote. One candidate aligns with the expert consensus on those nine issues while the other does not. There is expert disagreement about the last of the ten issues, which the citizen reasonably takes to be substantially less important than the nine about which there is consensus. An implausibly extreme interpretation of my central claim would maintain that these circumstances nullify the duty to vote. I do not want to make that implausibly strong claim. In such a case, one plausibly has a duty to vote for the candidate that aligns with expert consensus on the other issues.

But I do not want to defend the weak claim that when there is expert disagreement about *all* of the issues, the duty to vote is nullified. I rather am interested in an intermediate position: when there is widespread expert disagreement about *sufficiently many* issues central to the vote, the duty to

¹ Brennan (2011) and Maskivker (2019) hold that someone who votes on the basis on unjustified belief does not vote well.

vote is nullified. What counts as “sufficiently many” varies with facts of the case, e.g., the importance of the issue(s) about which there is expert disagreement.²

My primary contribution in this paper is to show how expert disagreement about political issues *can* have important implications for the duty to vote. To what extent such circumstances actually do obtain is an empirical issue that I won’t aim to fully settle here. But I suggest that these circumstances are not unrealistic or especially uncommon.

Readers may be tempted to think that such circumstances are unrealistic or uncommon if they focus on highly salient cases like national US general elections. In particular, one might think that there is one major issue that is subject to expert consensus and about which the two parties in US national elections disagree: climate change.³ In that case, voters can vote on the basis of this issue only.

But I urge readers to remember that there are many other sorts of votes, beyond US national elections, that take place; and that there are many hot-button political issues that are pervaded by expert disagreement. Here are three kinds of examples of votes that will often turn on such issues. First, many US states and other jurisdictions employ referenda, which typically center on a single policy question that may well be subject to expert disagreement. Second, before US general elections there

² For ease of exposition, I will use the phrase ‘circumstances of expert disagreement’ to mean ‘circumstances where there is expert disagreement about sufficiently many issues central to the vote.’

³ Many Republican officials profess skepticism about the reality of anthropogenic climate change. In this, they are out of sync with the expert consensus, which is that anthropogenic climate change is real. See So (2024).

Some reporting suggests that the Republican party is shifting from climate skepticism to a more nuanced position according to which they question whether the costs of particular responses to climate change are worth bearing (Jacobo, 2024) – a normative issue that is plausibly not within the purview of experts. Insofar as this is true, it only makes my argument here more pressing. The Republican party being out of sync with expert consensus about climate change is one path to preventing my argument from applying to the case of U.S. national politics, i.e. from nullifying the duty to vote in that case. If the Republican party is no longer out of sync with the expert consensus, that path is blocked.

are party primaries, which are often contested by candidates who are broadly aligned on the most important issues about which there is expert consensus (such as climate change) and differ only with respect to matters about which there is expert disagreement. Finally, and most obviously, many other countries are not like the US in that they do not have just two parties, one of which is broadly in line with the expert consensus on climate change and one that is radically out of line with it. In many other countries, it is common ground between parties that anthropogenic climate change is happening, and disputes about environmental policy concern the most appropriate response to this fact (which *is* subject to expert disagreement). Moreover, many other countries have multiparty systems wherein, even if one party is out of line with the scientific consensus, more than one party is broadly in line with it, and the important differences between *these* parties turn on issues subject to expert disagreement. As such, circumstances of expert disagreement are plausibly common enough that their implications for the duty to vote is not an issue that we can ignore.

2. Consequentialist Justifications of the Duty to Vote

Popular discourse tends to assume that citizens in democracies have a duty to vote. Organized campaigns and informal social pressure spring up around election time, pushing citizens to vote and admonishing them if they do not. But scholars of politics have been suspicious of this claim. Beginning with Anthony Downs (1957), scholars have noted that the probability of one's vote changing the outcome is vanishingly small. Consequently, the expected utility that voting will provide voters is quite low.⁴ Given that this is so, citizens are better off spending their time doing other things, according to standard decision-theoretic reasoning.

⁴ This is on the assumption that what determines the utility for the voter is the electoral outcome and not something like its function in expressing their political identity, something which Downs himself notes. My argument is ultimately going

Loren Lomasky and Geoffrey Brennan (2010) further develop the case against the duty to vote. The decision-theoretic reasoning introduced by Downs applies not only to the case where citizens are reasoning self-interestedly. Even if citizens are concerned about justice and the well-being of their compatriots, it *still* does not make sense, on consequentialist grounds, to vote. The expected utility of voting, even if we factor in the potential benefits to others, is still quite low because the probability of one's vote changing the outcome is so low. So even factoring in the ethical implications of voting, there does not seem to be a strong case in favor of it being obligatory.

Zach Barnett (2020) (building on an argument of Derek Parfit's [1984]) has pushed back against this line of reasoning. He argues that voting can be rational even when the probability of one's vote making a difference is very, very low because the utility associated with the outcome of the vote can be very, very large. Indeed, the probability of one's vote making a difference and the utility associated with the outcome of the vote seem to be inversely correlated, with the number of voters being a mediating variable. So, yes, when there are more voters, one's vote is less likely to make a difference; but this lowered probability is decision-theoretically canceled out by the greater utilities at stake when more people are affected by government action.⁵

More precisely, Barnett identifies two specific conditions that are jointly sufficient to make voting rational on consequentialist grounds. The first is the stakes condition: "The average social benefit of electing the better candidate is more than twice as great as the cost of voting" (427), where

to center on the *duty* to vote (as opposed to instrumental reasons to vote). The duty to vote is most plausibly grounded in how the electoral outcome affects others in one's polity. That so, what factors into the calculation of the expected utility for an *individual* is less relevant to my purposes.

⁵ I am assuming that population size and number of voters are roughly though imperfectly correlated.

the average social benefit is the average utility each person receives from the election of the better candidate as compared with the utility they receive from the election of the worse candidate. The second condition is the chances condition: “The probability of casting the deciding vote is at least one divided by the number of citizens” (427). Plausibly, Barnett thinks these conditions obtain in many realistic situations.

Julia Maskivker (2019) provides another kind of argument for the duty to vote. She provides a line of reasoning that can be described as *rule-consequentialist* (though this is not a term she invokes). She does not base the duty on the expected utility of the individual vote. Instead, she thinks the value of an individual vote is derivative on the value of the collective process to which it contributes. To *contribute*, in the relevant sense here, is not necessarily to make a causal difference in the sense that without one’s individual participation, the collective process would be any less valuable. Rather, it is something more like playing fair in doing one’s part, and not free riding on the contributions of others. The rationale for voting is thus not that voting might make a difference to the outcome, but rather that one is obligated to do one’s part in participating in a process that is, taken as a *whole*, momentarily valuable.

Maskivker takes the fact that the process of voting is so immensely valuable and pairs it with the further premise that voting is relatively low cost. The thought is that when something contributes to an immensely valuable process and is relatively easy to do, you are morally obligated to do it.⁶

Before moving on, I will address an alternative consequentialist-style argument in favor of the duty to vote. Perhaps voting can contribute to a good outcome that is independent of the *electoral outcome*. That is, perhaps voting produces good consequences regardless of who one votes for and

⁶ See also Goldman (1999).

regardless of what policy/candidate ultimately wins. The relevant result-independent consequences are the perception of the state's legitimacy. If more people voting means citizens perceive the state as more legitimate, and people perceiving the state as more legitimate affords the democracy greater stability, then perhaps voting is a good contribution, whatever the outcome.

This line of thought depends on the empirical premise that higher turnout increases perceptions of legitimacy. While there is much empirical work on turnout, its focus is often on other relationships, e.g. whether higher turnout would affect demographic representation (Lijphart, 1997; Lutz & Marsh, 2007) or to what degree perception of legitimacy affects turnout (Blondel, Sinnott, & Svensson 1998, chapter 3). One exception is Kirkland & Wood (2017). They find that perceptions of legitimacy are not straightforwardly linked to turnout. Low turnout may or may not be seen as an indicator of illegitimacy, depending on what elite actors say about it. While this does not altogether undermine any relationship between turnout and perceived legitimacy, it significantly complicates it. It means that sometimes greater turnout may come with the electoral-outcome-independent good of perceived legitimacy, but sometimes it may not. Insofar as this is so, any argument for the duty to vote based on these sorts of consequences is restricted in scope.

Thus, I think the consequentialist justification that appeals to the consequences of the *electoral outcome* is the more promising path to a duty to vote. That more promising path has its own issues, though, as I will begin to demonstrate in the next section.

3. The Cost of Voting

Barnett's and Maskivker's arguments share an important premise: voting is not particularly costly. This premise is explicit in Maskivker's argument. It is less explicit in Barnett's argument, but still necessary for the argument to go through. Remember that voting is only act-consequentially

justified, on Barnett's picture, when the average social benefit of electing a given candidate is twice the cost of voting. Part of what makes it plausible that that this condition is often met is the thought that the cost of voting is pretty low. It requires only a couple of hours to go to the polling center.

Both arguments, then, rely on the claim that voting is (1) low-cost. But they also rely on the claim that it is (2) a way to achieve (or contribute to) a large amount of good. Unfortunately, (1) and (2) are in tension. Voting is only obviously low-cost if we consider only the act of going to the polling center and casting a ballot. But achieving the requisite large amount of good requires more than this. It requires not only that one vote, but that one vote *well*. Maskivker is explicit about this: she adds to the outright duty to vote a conditional duty to vote on the basis of justified beliefs—if one does vote, one ought to vote on the basis of justified beliefs. But inquiry enough to form justified beliefs about social and political issues can be quite time-consuming. It is not obvious, without further argument, that doing so counts as relatively low-cost.

We can identify a similar issue in Barnett's argument. The average social benefit of electing the candidate one votes for is twice that of the cost of voting *only if* one in fact votes for the better candidate. But doing this reliably plausibly requires forming justified beliefs about the campaign issues so that one can identify the better candidate. Again, it is not only the cost of going to the polling center and casting a ballot that is relevant; the cost of all the legwork needed to identify the better candidate is relevant as well. With this complication added, Barnett's sketch of why we should expect his stakes condition to be met is insufficient.

Perhaps, even with the costs of forming justified beliefs included, a cost-benefit analysis of voting will still support the duty to vote. Whether this is so depends on complex empirical issues. Brennan and Freiman (2022) argue convincingly that, in general, that it will not. The complexity of the issues that distinguish candidates makes it exceedingly difficult to form justified beliefs about

which candidate will yield better outcomes. My argument in this paper is in line with those claims of Brennan and Freiman, but importantly tackles the issue from an alternative angle in order to make an even more decisive claim. While Brennan and Freiman do an admirable job in showing that it will often be quite difficult to form justified beliefs about the issues central to a vote, I am going to show that there are realistic circumstances in which this is simply *impossible* to do—namely, circumstances of expert disagreement. If this is so, then under circumstances of expert disagreement, there is no duty to vote.

Let me briefly address one lingering issue. I have so far been proceeding in essentially consequentialist terms. But one might wonder whether there are cases where non-consequentialist concerns over respect for basic rights take precedence over concerns that turn on the empirical consequences of a policy (and where the expert disagreement concerns the latter). It is possible for there to be cases like this, but I want to note that it will often be far from straightforward. Even when non-consequentialist values are at stake, empirical matters will still likely be relevant. First, empirical facts can determine whether a right is in fact violated. Second, most non-consequentialists still think that consequences matter morally, even if they don't think they are all that matters morally. In both cases, there might be expert disagreement over the relevant empirical facts. So, while non-consequentialist values may sometimes play a role in establishing a duty to vote or preventing its nullification, things are certainly not so simple as to say that whenever non-consequentialist values are in play, the issue posed by expert disagreement no longer applies.

4. The Epistemology of Expert Disagreement

There is much work in epistemology concerning whether laypeople ought to *defer* to experts. Experts have cognitive advantages over laypeople. They have the requisite background knowledge, formal training, and experience in the relevant field. Thus, laypeople plausibly ought to defer to experts

about the issues on which they have expertise.⁷ On the other hand, some think that we can go too far in deferring to experts. Trusting in experts too much can make us vulnerable to be misled.⁸

My concern here is slightly different. I am focused on what laypeople ought to do when the experts *disagree* with one another. In particular, I am concerned with circumstances of *pervasive* expert disagreement. There is some vagueness in this characterization, owing to the term ‘pervasive.’ When ninety-nine percent of experts agree and only one percent dissents, that is clearly not pervasive disagreement. And when fifty percent of the experts think one thing and fifty percent think another, that clearly is pervasive disagreement. There is some vagueness between those two extremes. I do not think the experts have to be split exactly down the middle for their disagreement to count as pervasive. Fifty-three percent of experts asserting p and forty-seven percent asserting not-p seems to me to count as the sort of pervasive expert disagreement that will have the epistemic implications I discuss. Whether a seventy-thirty split counts is less clear. This vagueness is fine for my purposes; I leave it up to others to determine whether instances of expert disagreement count as pervasive on a case-by-case basis.

A central premise in my argument is that when there is pervasive expert disagreement about some topic, laypeople ought to suspend judgment about that topic. I take suspension of judgment to be a coarse-grained doxastic attitude, a third option in addition to belief and disbelief. As Scott Sturgeon (2020) puts it, it is an attitude of “committed neutrality.”

⁷ Thomas Grundmann (2021), for example, has defended a particularly strong version of this claim.

⁸ See Hazlett (2016) for an argument that refraining from deferring to experts can be socially valuable. See Lackey (2021) for a concern that deferring to experts too uncritically can make one liable to manipulation.

The claim that laypeople ought to suspend judgment under conditions of expert disagreement is defended by a number of philosophers. Michael Huemer (2005) writes: “suppose the issue is one about which there is no consensus among the experts... In this case, [suspension of judgment] is advisable” (p. 525). Elizabeth Anderson (2011) writes: “When the vast majority of diverse inquirers converge on certain conclusions... a robust scientific consensus obtains. Before a consensus, the best course for laypersons is to suspend judgment” (p. 149).

Widespread disagreement indicates two things. First, that the testimonial evidence from the experts – the evidence we get from the experts reporting their views – is inconclusive. If a substantial number of experts testify that p , and a substantial number testify that $\sim p$, then the testimonial evidence that we laypeople have is itself inconclusive. Thus, we cannot come to justifiably form a belief about p on the basis of the testimonial evidence from the experts.

But what if we can appeal to some other non-testimonial evidence in order to tip the scales and come to justifiably believe one way or another? This leads to the second lesson of the existence of expert disagreement: the scientific evidence on which the experts base their views is likely either inconclusive or highly complicated. If it weren't, then it's unlikely that there would be widespread expert disagreement about the issue. The fact that the experts disagree suggests that the evidence is difficult to work through and does not straightforwardly support any clear interpretation. But this means that our own hopes, as laypeople, of working through that more direct scientific evidence are even worse. If the evidence is complicated for those who have training and experience, then it's highly unlikely that we laypeople would be able to reliably evaluate it.

When we combine these two lessons, we see why laypeople ought to suspend judgment under circumstances of expert disagreement. The evidence that is accessible to us, the testimony from the experts, is inconclusive. As for the more direct scientific evidence, even if it's sufficient to justify a

belief, it's likely inaccessible to us. When the experts disagree, we can't learn from them, and we can't make headway ourselves; we ought to suspend judgment.

Before moving onto the next section, where I apply insights from this section to the debate about the duty to vote, I want to note a response one might consider. If 55% of experts think that p and 45% think that $\text{not-}p$, why do such conditions necessitate suspension of judgment? Why can't we believe what the slight-to-moderate majority believes? While this may seem promising, I think we ought to heed Goldman's (2001) warning against solely 'going by the numbers,' particularly in cases of slight-to-moderate majorities. What matters is not the bare number of experts who judge that p , but rather the number of experts that *independently* judge that p .⁹ That there is a slight-to-moderate majority doesn't put us in a good position to infer that the majority has a sufficient degree of judgment independence. So a slight-to-moderate majority cannot epistemically justify us in forming a belief about the issue.

One might have a related thought here: even if we must suspend judgment about the issues in question, might we still rationally take action, despite our having suspended judgment? I will address this important line of thought in detail in section 6.1.

5. Expert Disagreement and the Duty to Vote Well

We now have our key ingredients in hand. Voting well requires having justified beliefs about the issue one is voting on. Laypeople ought to suspend judgment about issues on which the experts disagree. This is where our tension arises. A concrete case will illustrate.

⁹ There is some controversy about Goldman's arguments (Coady, 2006; Lackey, 2013). Space restrictions prevent me from wading into this debate, but I will note that I think Goldman's arguments are most compelling in cases of slight-to-moderate majorities. The problems for the argument primarily arise in cases of large majorities or consensus (something that Goldman acknowledges).

In 2015, the Chicago Booth's Initiative for Global Markets polled 43 professional economists, asking them to state their degree of agreement or disagreement with the following statement: "If the federal minimum wage is raised gradually to \$15-per-hour by 2020, the employment rate for low-wage US workers will be substantially lower than it would be under the status quo." 26% of respondents indicated that they strongly agreed (5%) or agreed (21%) with the statement; 24% of respondents indicated that they disagreed with the statement; 38% indicated that they were uncertain.¹⁰ This is, to my mind, a clear example of expert disagreement. A year later, citizens of the state of Washington voted on Initiative 1433, a ballot to increase the state's minimum wage from \$9.47 to \$13.50 by 2020.

Given the expert disagreement about the minimum wage, Washington voters could not have had justified beliefs about the outcome of policy they voted for. So they couldn't have voted on the basis of justified beliefs. And, of course, if they refrained from voting, they wouldn't have voted well either, because they wouldn't have voted. Thus, under these circumstances of expert disagreement, it was impossible for Washington voters to discharge any purported democratic duty to vote well. Assuming that ought implies can, this implies that in fact, they were under no such duty.

Thus goes the argument against the outright duty to vote well. Given that the outright duty to vote and the conditional duty to vote well entail the outright duty to vote well, we must also give up one of the two former duties (at least in full generality). Which one should we give up? There are two reasons to think that we ought to give up the outright duty to vote well rather than the conditional duty to vote well.

The first reason is that there is a compelling argument in favor of the existence of the conditional duty to vote well. Indeed, both proponents of the outright duty to vote, like Maskivker,

¹⁰ <https://www.igmchicago.org/surveys/15-minimum-wage/>.

and skeptics, like Brennan (2010), accept the existence of the conditional duty to vote well. Brennan argues that voters without justified beliefs vote either harmfully or fortuitously. The former “occurs when a person votes, without epistemic justification, for harmful policies or for candidates likely to enact harmful policies,” whereas the latter “occurs when citizens vote for what are in fact beneficial policies or candidates likely to enact beneficial policies, but they lack sufficient justification to believe that these policies or candidates are good” (p. 68). Even in the latter case, Brennan argues, voters impose an unjustified *risk* on their fellow citizens. Imposing risks can be unjustified even if no harm comes about. While it is sometimes morally permissible to impose risks on others, one must do so on the basis of justified beliefs. It’s because voting on the basis of unjustified beliefs is either harmful or fortuitous, Brennan argues, that we ought not do it—hence establishing the conditional duty to vote well.

The second reason to give up the outright duty to vote rather than the conditional duty to vote well is that the best arguments for the outright duty to vote presuppose the conditional duty to vote well. The reason we have a duty to vote, according to those arguments, is that voting can achieve some good. But it can only (reliably) do so if it’s done well. Brennan and Lomasky (2010) note this in their extended critique of potential arguments for the duty to vote: the prominent arguments in favor of the duty only follow if we assume voters are voting well. Hence, we can’t give up the conditional duty to vote without giving up the outright duty to vote as well. With this being so, the way to go is to conclude that there is no outright duty to vote (at least in circumstances of expert disagreement).

With the whole argument now completed, I will now address some objections to its crucial step, the argument that the outright duty to vote well is impossible to discharge under circumstances of expert disagreement. This will involve looking at some alternative characterizations of the epistemic conditions on voting well.

6. Alternative Epistemic Conditions on Voting Well

Following Brennan and Maskivker, I have characterized the epistemic condition on voting well in terms of having justified *beliefs*. But a proponent of the duty to vote might suggest that the epistemic condition on voting well should be characterized differently. First, they might suggest we can make do solely with justified credences, rather than justified beliefs. Second, they might appeal to the miracle of aggregation or the Condorcet Jury Theorem. I argue that neither of these alternative characterizations of voting well assuages the problems for the duty to vote raised by expert disagreement.

6.1 *Credences and Expected Utility*

First, let's consider the proposal that we should characterize the epistemic condition on voting well in terms of *credence*. On this view, voters do not have to have outright beliefs about the effects of the policies or candidates they vote for in order to vote well; it is enough that they have (justified?) credences about these effects. The import of this proposal becomes most clear if we look to some basic decision theory.

Consider a voter in the Washington minimum wage referendum. Suppose they have learned that the experts disagree about the effect of minimum wage on unemployment and, being epistemically rational, they have suspended judgment about that relationship. But they still have some level of confidence in the proposition that raising the minimum wage increases unemployment. They also have some sense of how bad the outcome would be if unemployment rose. They thus have what they need to calculate the expected utility of raising the minimum wage. According to basic decision theory, this is enough to allow them to rationally decide between their voting options.

Voting on the basis of the expected utility of the policy or candidate one votes for—where this expected utility is informed by one's (justified) credences, not outright beliefs—does seem to

count as an instance of voting well. It does not seem to impose unjustified harms or risks on one's compatriots. Moreover, I have argued only that expert disagreement requires laypeople to suspend judgment; my arguments have said nothing about credences. If all this is correct, then my argument against the outright duty to vote well falls flat.

However, I think this critique goes wrong in suggesting that credences entirely remove the need for belief in calculations of expected utility. Consider this illustration from Brian Weatherson (2012):

Professor Dec is teaching introductory decision theory to her undergraduate class. She is trying to introduce the notion of a dominant choice. So she introduces the following problem, with two states, S1 and S2, and two choices, C1 and C2, as is normal for introductory problems.

	S1	S2
C1	-\$200	\$1000
C2	-\$100	\$1500

...One student, Stu, is particularly unwilling to accept that C2 is better than C1. He thinks, on the basis of his experience, that when more than \$1000 is on the line, people aren't as reliable about paying out on bets. So while C1 is guaranteed to deliver \$1000 if S2, if the agent bets on C2, she might face some difficulty in collecting on her money.

Weatherson says that Stu is going wrong somewhere, but he notes that it is not obvious how. What is Stu's error? According to Weatherson, it is roughly that he fails to recognize that every outcome placed on the decision matrix is a *known* outcome. We *know* that if we place bet C2 and the world is in state S2, then we get \$1500. That is why Stu's complaint is misguided. This leads Weatherson to claim,

generally, that in order to fill in a cell of a decision matrix with an outcome, we must know that that outcome obtains in that act-state pair.

But let's weaken Weatherson's constraint a bit. I think the core of the issue is that an outright *belief*, rather than a credence, is necessary in order to fill in the cells of a decision matrix.¹¹ It will be helpful to see why this is if we consider Stu's complaint in a context in which we *don't* believe that bet C2 will yield \$1500 in S2. In such a case, we cannot simply dismiss Stu's protests. Rather, what we ought to do is update our decision matrix to something like the following:

	S1	S2	S3
C1	-\$200	\$1000	\$1000
C2	-\$100	\$1500	\$0

where S2, in this context, is the state where C2 wins and the people paying out are trustworthy and reliable, and S3 is the state where C2 wins and people are not trustworthy and reliable. (Suppose, for instance, that the payer in S3 will accept their fate when they owe \$1000; they will pay up. \$1000 is a lot, but not enough for them to go into hiding. But \$1500 is enough money for it to be worth it for them to go into hiding, in which case we won't be able to collect.) What seems to be going on is that Stu assigns some amount of credence to S3. This decision matrix can accommodate that.

¹¹ Savage (1954) and Jeffrey (1965) both suggest this as well. They contend that in setting up our decision matrices, we have to treat some things as if we have outright beliefs in them, setting aside any grounds for uncertainty.

But suppose Stu is feeling particularly obstinate. Now he also thinks that people won't entirely shirk payouts of \$1000, but that they will sometimes shortchange you by \$50. Again, the decision-theoretic approach can accommodate this, by complicating things still further:

	S1	S2	S3	S4	S5
C1	-\$200	\$1000	\$1000	\$950	\$950
C2	-\$100	\$1500	\$0	\$1450	\$0

But suppose Stu *still* isn't satisfied. Suppose he thinks some people are graceful in defeat, and they'll throw in an extra \$50. We can see where this is going. We can keep complicating the decision matrix to accommodate Stu's concerns, but the buck has to stop somewhere. Eventually, if Stu is to calculate the expected utility of his options, Stu has to settle on a matrix such that he has a belief, for each act-pair CN-SN, that some outcome will result. Otherwise a regress looms.

With this in mind, let's return to the case of the Washington minimum wage referendum.

Here's the start of the decision matrix the voters will have to fill out:

	Minimum wage \rightarrow unemployment	\sim (Minimum wage \rightarrow unemployment)
Increase minimum wage		
Don't increase minimum wage		

What did we learn from the discussion of Weatherson? In order to fill in those cells, voters need to *believe* what will happen in each act-state pair. They need to have a belief, for instance, about what will happen if the state increases the minimum wage and minimum wage increases unemployment, *just as Stu had to have a belief about what would happen were he to place a certain bet in a given state*. Stu's outright beliefs about the outcomes in each choice-state pair are what allowed him to fill in his decision matrix and thus calculate expected utility. Similarly, it is our voters' outright beliefs about the outcomes in each choice-state pair that allow them to fill in decision matrices and calculate expected utility.

Why is this important? Well, it raises our problem of expert disagreement all over again. The experts could very well be in disagreement not just about whether a rise in the minimum wage will increase unemployment, but about what will happen if unemployment increases. If that's so, then laypeople ought to suspend judgment about what will happen if unemployment increases, i.e. they do not have any belief. Then they cannot fill in the corresponding cell of the decision matrix, so they cannot calculate the expected utility of increasing the minimum wage, which was supposed to be the new basis of their vote.

Of course, they could do the same thing as our obstinate version of Stu. They could complicate the decision matrix further and further to accommodate all the expert disagreements. That is, they could locate the source of the experts' disagreement about the relationship between minimum wage and unemployment and adjust the decision matrix accordingly.

If the experts disagree about what happens when we raise minimum wage, what is the basis of that disagreement? Suppose it is the effect of unemployment on inflation. Some experts think unemployment eases inflation, and some think it does not. With this so, voters can construct the following decision matrix:

	Minimum wage → unemployment & unemployment → lower inflation	Minimum wage → unemployment & ~(unemployment → lower inflation)	~(Minimum wage → unemployment) & unemployment → lower inflation	~(Minimum wage → unemployment) & ~(unemployment → lower inflation)
Increase minimum wage				
Don't increase minimum wage				

Now the propositions about the effects of unemployment are in effect moved from the cells, where voters need but cannot justifiably have beliefs about them, to the columns, where they need only credences in them. But you can see, looking to our example of obstinate Stu above, where this is heading. Suppose there is further expert disagreement about what outcome will occur as a result of inflation lowering. Perhaps that outcome depends on a change in some other economic metric. If this is so, then we have to complicate the decision matrix even further. Each time we do so, we add more cells, which means we add more propositions about which there cannot be expert disagreement. If any one of the cells is the subject of expert disagreement, we cannot calculate the expected utility of at least one of the choices, and thus we cannot compare the expected utility of the two.

Just like Stu, prospective voters eventually have to settle somewhere, with beliefs about which outcomes obtain in each act-state pair—on pain of regress. But if they cannot have justified beliefs in

propositions under conditions of expert disagreement, this means that, to be responsible, they have to set up the matrix such that the propositions about what outcomes eventuate in each of the cells are not subject to expert disagreement. But now the question for the (credence-friendly) proponent of the duty to vote is: why expect that there *will* be a setup of the matrix such that this is so? In fact, if anything, the more complicated the matrix becomes, the *less* likely it is that the propositions about what outcomes eventuate in each of the cells are not subject to expert disagreement. For the more complex the matrix, the more propositions about which outcomes eventuate there are, and the more likely it is that at least some of them are subject to expert disagreement. Complicating the matrix provides *more* opportunities for expert disagreement to interfere with voters' ability to vote well, not fewer.

I don't claim that it will *never* be possible for a voter to set up a decision matrix (about some contested issue) such that the propositions about what outcomes eventuate in each of the cells are not subject to expert disagreement. In order for it to be possible, we would need what we can call an instance of *isolated* expert disagreement, where experts in some domain disagree about relatively few questions/issues in that domain and have achieved consensus on the others. In such cases, voting on the basis of expected utility might circumvent the problem of expert disagreement. But we have no reason to expect expert disagreements in general to be isolated, and we certainly don't have reason to think that *all* expert disagreements are isolated. Thus, at most, the shift to the credence-based, expected-utility framework makes my argument apply in slightly fewer cases.

It might now be objected that, to vote responsibly, I do not even need to be able to rationally think that I am voting for the (expectedly) best option. For if I *don't* vote, others still will, and those people might be even *less* able to identify the (expectedly) best option than I am. So perhaps all I need

to think is that I am more likely to vote for the expectedly best option than others. Perhaps that suffices for the expected utility of my voting to be positive, and thus for my having the duty to vote.

I think this is an important complication, but I don't think it goes far enough to save the duty to vote in circumstances of expert disagreement. For one, this consideration would only apply to those who can rationally think that they are epistemically better than others. This is far from establishing a general duty to vote. Moreover, people notoriously overestimate their epistemic capabilities relative to others¹², and so many people who *think* that they are epistemically better than others may not be rational in so thinking. Since we should be wary of this danger, we should be very cautious about voting on the basis of the thought that we are epistemically better than others.

6.2 The Miracle of Aggregation

Another characterization of the epistemic condition on voting appeals to the 'miracle of aggregation.' According to the miracle of aggregation, it is not the case that participants in collective decision-making need to be individually reliable or justified in order for that procedure to reliably produce correct outcomes. Roughly put, the miracle of aggregation is the phenomenon by which the errors made by individuals cancel each other out in collective decision making.¹³

The guiding thought here is that voting well does not require being individually justified or individually reliable. Rather, it requires only that one contribute to a procedure which is ultimately reliable. Individual voters, even if they are mistaken or unjustified, play an important role in the reliability of the aggregation process. Their errors cancel out the opposing errors of others. And so

¹² See Alicke & Govorun (2005) for an overview of the better-than-average effect, the social-psychological tendency to rate oneself as better than average across a wide range of domains. Pronin, et al. (2002) finds specifically that people view others as being more cognitively biased than themselves. Zell et al. (2020) conduct a meta-analysis that confirms the robustness of the better-than-average effect.

¹³ For a discussion of the miracle of aggregation and its relevance to democratic politics, see Landemore (2012).

their participation is required, and that requirement is grounded in the consequences of the decision making, as I have argued is necessary to justify the duty to vote.

There are several different versions of the miracle of aggregation, which appeal to different explanations of how aggregation is supposed to yield greater reliability than the individuals that comprise it. The miracle has its critics.¹⁴ I will not re-hash their critiques here. Rather, I will note various assumptions that seem to be required by each version of the miracle of aggregation and suggest why the existence of expert disagreement ought to make us skeptical of those assumptions.

On the elitist explanation of the miracle of aggregation, there is an elite core of participants in the collective decision making that are in fact reliable. If the errors of those outside that elite core are not correlated, then if the group is sufficiently large, those errors will cancel each other out. With the errors being cancelled out, the input of the elite core wins the day. An important assumption of the elitist version is that there is some elite core of reliable participants, i.e., the experts. But the existence of widespread expert disagreement casts some doubt than any such elite core is in fact reliable.

On the democratic explanation of the phenomenon, everyone has got matters roughly correct, with a bit of noisy error around their judgments. As long as the noisy error surrounding individuals' judgments is not correlated, then those errors will cancel each other out. The democratic explanation assumes that participants have got the matters relevant to the vote roughly correct. But recall that expert disagreement suggests that the topic at hand is a complicated one; assessing the relevant evidence is a demanding task, one perhaps beyond the ken of laypeople. So the existence of expert disagreement casts some doubt on the assumption that laypeople have things roughly correct.

On the distributed explanation of the phenomenon, individuals in collective decision making have various bits of relevant information that they bring to bear on the question at hand. Those bits

¹⁴ See Caplan (2008, 2009) for overviews of concerns about the applicability of the miracle of aggregation.

of relevant information come together via aggregation to reliably yield a correct outcome. But, for reasons similar to those outlined in the previous paragraph, the existence of expert disagreement should make us hesitant to assume that the relevant information is distributed among participants in collective decision making—or, at least, that they are able to reliably evaluate that distributed information. Expert disagreement suggests that the topic is too complex for laypeople to make sense of.

On any understanding of the miracle of aggregation, then, the existence of widespread expert disagreement should raise some suspicion that the necessary assumptions for the miracle obtain. Nothing I have said constitutes a knockdown defeat of the miracle of aggregation's assumptions. My aim has been more modest because I think the burden I face is more modest. I have tried to raise some plausible considerations that should make us question the plausibility of the miracle of aggregation's various assumptions obtaining under conditions of expert disagreement. But even if these are not knockdown proofs that the assumptions do obtain, we also lack knockdown reasons to think that the assumptions of the miracle do obtain. It's a wash.

6.3 The Condorcet Jury Theorem

A final attempt to recharacterize the epistemic condition on voting well appeals to the Condorcet Jury Theorem (CJT). According to the CJT, provided that participants (1) vote independently, and (2) are more than 0.5 reliable with respect to the (binary) question at issue, then, as the number of participants increases, the probability that the vote yields the correct outcome approaches 1.¹⁵ With this being so, one might think that it is inappropriate to characterize voting well

¹⁵ For discussion of the Condorcet Jury Theorem and its application to democratic decision-making, see Anderson (2006), Estlund (2009), and Landemore (2012). Critics often contend that the CJT is inapplicable to democratic decision-making because voters do not in fact vote independently. For a discussion of and response to this objection, see Dietrich & Spiekermann (2013). I will focus on a different line of critique.

in terms of voters' justified beliefs. As long as voters are greater than 0.5 reliable with respect to some question, then it doesn't matter what their epistemic status of their beliefs is. If they all contribute (independently), then the procedure is very likely to converge on the correct outcome.

But this doesn't take seriously the force of Brennan's risk-based arguments against voting without justified beliefs. That argument contended that even if a citizen's vote ended up contributing to a beneficial outcome for their compatriots, nonetheless that vote was morally problematic. Even though things turned out alright, it imposed an unjustified risk on their fellow citizens. Something similar can be said here. Even if a citizen just so happens to be greater than 0.5 reliable with respect to the issue they are voting on, unless they have justified beliefs about the outcome, they are imposing an unjustified risk on their compatriots.

But what if it's not merely that citizens *just so happen* to be greater than 0.5 reliable with respect to the issue being voted on? What if they *justifiably believe* that they are greater than 0.5 reliable? If that were so, it is much less convincing to say that they are imposing unjustified risks on their fellow citizens. But it's not clear that citizens can justifiably believe they are greater than 0.5 reliable under conditions of expert disagreement. The experts' reliability is in question in such cases; the laypeople's even more so. The fact of expert disagreement makes laypeople's belief in their own reliability, even greater than 0.5 reliability, much more dubious. If I were to give you a true or false question on some complicated topic and inform you that experts in the relevant field have all been providing different answers, it doesn't seem justified for you to think your guess is going to be any better than chance.

Proponents of the CJT tend not to conduct empirical studies to determine whether voters are in fact 0.5 reliable with respect to the issue being voted upon. Indeed, what indirect empirical evidence we have suggests the opposite. Caplan (2008), Brennan (2011 & 2016), and Somin (2016) all

extensively document the extent of voter ignorance, especially among American voters. Voters are often unaware of basic governmental procedures and the identities of those who represent them, let alone the details of the complex issues that are at stake in deciding between policies and candidates. These are all *prima facie* reasons to doubt that the reliability assumption obtains. What I want to note here is that the existence of expert disagreement ought to be added to the list of *prima facie* reasons to doubt this assumption. This is not to say that expert disagreement definitively entails that voters are less than 0.5 reliable. It only undermines our entitlement to that assumption. But, again, the proponent of the CJT did not definitively show that voters are greater than 0.5 reliable. The burden is on the proponent of the CJT to show that the assumption is met in response to *prima facie* defeaters to it.

7. Conclusion

I have tried to show that under certain realistic circumstances—circumstances of expert disagreement—one cannot vote well, on a variety of epistemic characterizations of what voting well entails. Expert disagreement prevents us from formulating the justified beliefs we need to vote well, and there is nothing we can do about it – we cannot make the experts come to a consensus. It thus appears that, when an instance of democratic decision-making centers on sufficiently many issues that are subject to expert disagreement, citizens do not have a duty to vote.

As discussed in the introduction, one might think that the duty to vote is not nullified in US national elections, given the parties' disagreement about climate change. I want to conclude by noting a certain kind of perverse fragility to this way of avoiding the problem of expert disagreement. Ideally, when there is genuine expert consensus, political candidates will conform their policies to it. Suppose this happens; the party out-of-sync with the consensus changes their tune and aligns with the expert consensus on climate change. Good news! But yet, now the problem of expert disagreement re-arises. Then, what distinguishes the two parties are matters about which there is expert disagreement, so

citizens have no basis on which they can justifiably vote. One would have thought that representatives aligning with expert consensus was an unabated good! But in fact it here raises new problems. This bizarre implication captures the importance of expert disagreement for democratic politics.

Bibliography

\$15 minimum wage. (2015). Kent A. Clark Center for Global Markets.

<https://www.kentclarkcenter.org/surveys/15-minimum-wage/>

Alicke, Mark D., & Govorun, Olesya (2005). The better-than-average effect. In M. D. Alicke, D. A. Dunning, & J. I. Krueger (Eds.), *The Self in Social Judgment* (pp. 85–106). Psychology Press.

Anderson, Elizabeth (2011). Democracy, public policy, and lay assessments of scientific testimony. *Episteme*, 8(2), 144–164.

Anderson, Elizabeth (2006). “The epistemology of democracy.” *Episteme* (1-2):8-22.

Barnett, Zach (2020). “Why you should vote to change the outcome.” *Philosophy and Public Affairs* 48 (4):422-446.

Beerbohm, Eric (2012). *In our name: The ethics of democracy*. Princeton University Press: Princeton, NJ.

Blondel, Jean, Sinnott, Richard, & Svensson, Pale (1998). *People and parliament in the European Union: Participation, democracy, and legitimacy*. Clarendon Press ; Oxford University Press.

Brennan, Jason (2011). *The ethics of voting*. Princeton University Press: Princeton, NJ.

Brennan, Jason (2016). *Against democracy*. Princeton University Press: Princeton, NJ.

Brennan, Jason, & Freiman, Christopher (2023). Why swing-state voting is not effective altruism:

The bad news about the good news about voting*. *Journal of Political Philosophy*, 31(1), 60–79.

Caplan, Bryan (2008). *The myth of the rational voter: why democracies choose bad policies*. Princeton University Press.

- Caplan, Bryan (2009). Majorities against utility: implications of the failure of the miracle of aggregation. *Social Philosophy and Policy*, 26(1), 198–211.
- Coady, David (2006). When experts disagree. *Episteme*, 3(1–2), 68–79.
- Dietrich, Franz & Kai Spiekermann (2013). ‘Epistemic democracy with defensible premises.’ *Economics and Philosophy* 29 (1): 87-120.
- Estlund, David (2009). *Democratic authority: A philosophical framework*. Princeton University Press: Princeton, NJ.
- Goldman, Alvin (1999). ‘Why citizens should vote: A causal responsibility approach.’ *Social Philosophy and Policy* 16 (2): 201-17.
- Goldman, Alvin (2001). Experts: Which ones should you trust? *Philosophy and Phenomenological Research*, 63(1), 85–110.
- Grundmann, Thomas (2021). ‘Facing epistemic authorities: Where democratic ideals and critical thinking mislead cognition.’ In Sven Bernecker, Amy Flowerree & Thomas Grundmann (eds.) *The epistemology of fake news*. Oxford University Press: New York, NY.
- Hazlett, Allan (2016). ‘The social value of non-deferential belief.’ *Australasian Journal of Philosophy* 94 (10): 131-151.
- Huemer, Michael (2005). ‘Is critical thinking epistemically responsible?’ *Metaphilosophy* 36 (4): 522-531.
- Jacobo, Julia (2024, August 27). The GOP has transitioned from climate denial to climate misrepresentation, experts say. *ABC News*. <https://abcnews.go.com/Politics/gop-transitioned-climate-denial-climate-misrepresentation-experts/story?id=113056571>
- Jeffrey, Richard (1965). *The logic of decision*. University of Chicago Press: New York, NY.

- Kirkland, Christopher & Wood, Matthew (2017). Legitimacy and legitimization in low turnout ballots. *Government and Opposition*, 52(3), 511–531.
- Lackey, Jennifer (2013). Disagreement and belief dependence: Why numbers matter. In D. Christensen & J. Lackey (Eds.), *The epistemology of disagreement: New essays* (pp. 243–268). Oxford university press.
- Lackey, Jennifer (2021). 'Preemption and the problem of the predatory expert.' *Philosophical Topics* 49 (2): 133-150.
- Landemore, Helene (2012). *Democratic reason: Politics, collective intelligence, and the rule of the many*. Princeton University Press: Princeton, NJ.
- Lijphart, Arend (1997). Unequal participation: Democracy's unresolved dilemma. *American Political Science Review*, 91(1), 1–14.
- Lomasky, Loren & Geoffrey Brennan (2000). 'Is there a duty to vote?' *Social Philosophy and Policy* 17 (1):62.
- Lutz, Georg & Marsh, Michael (2007). Introduction: Consequences of low turnout. *Electoral Studies*, 26(3), 539–547.
- Maskivker, Julia (2019). *The duty to vote*. Oxford University Press: New York, NY.
- Parfit, Derek (1984). *Reasons and persons*. Oxford University Press: New York, NY.
- Pronin, Emily, Lin, Daniel Y., & Ross, Lee (2002). The bias blind spot: perceptions of bias in self versus others. *Personality and Social Psychology Bulletin*, 28(3), 369–381.
- Savage, Leonard (1954). *The foundations of statistics*. Wiley Publications in Statistics: New York, NY.
- So, Kat (2024). *Climate deniers of the 118th Congress*. Center for American Progress.
<https://www.americanprogress.org/article/climate-deniers-of-the-118th-congress/>

Somin, Ilya (2016). *Democracy and political ignorance: Why smaller government is smarter*. Stanford Law Books.

Sturgeon, Scott (2020). *The rational mind*. Oxford University Press.

Weatherson, Brian (2012). 'Knowledge, bets, and interests.' In Jessica Brown & Mikkel Gerken *Knowledge ascriptions*. Oxford University Press: New York, NY. pp. 75-103.

Zell, Ethan., Strickhouser, Jason E., Sedikides, Constantine, & Alicke, Mark D. (2020). The better-than-average effect in comparative self-evaluation: A comprehensive review and meta-analysis. *Psychological Bulletin*.