

# Evidence, Causality, and Collective Action

## 1. Introduction

If everyone votes for the better candidate, then that candidate will win. If no one does so, then the worse candidate will triumph, or the political system may even collapse (if no one votes). Similarly (to take another favorite example of philosophers'), if everyone refrains from buying factory-farmed meat, animal suffering will be greatly reduced, and gustatory pleasure will be (at most) somewhat lower than if everyone buys factory-farmed meat.

Given the good consequences of everyone acting in these ways, one might think that there is a compelling consequentialist rationale for each of us to do our part. However, the realization of these kinds of goods presents a collective action problem, which arises from the “lumpy” way in which the goods are realized. That is, there are critical zones at which every contribution makes a difference, but outside of these zones, each contribution makes no difference to the realization of the good in question.<sup>1</sup> Factory farms do not reduce production every time a consumer reduces his consumption of meat. So in many cases, a consumer’s choice will have no effect on animal suffering. Under majority rule, the electoral outcome will be the same, regardless of whether or how any given individual votes, unless the better candidate wins by a single vote.

Many consequentialists argue that even if any given contribution won’t *in fact* have good consequences in lots of these cases, it will still often have good *expected* consequences (Singer 1980; Parfit 1984; Norcross 2004; Kagan 2011; McMullen and Halteman 2018; Barnett; 2020; Hedden 2020; Isaacs, Lerner, and Russell forthcoming). Call this the *decision-theoretic response* to collective action problems. These theorists stress individuals’ ignorance about whether the total

---

<sup>1</sup> The term “lumpy goods” is commonly used by economists and political scientists (see, e.g. Taylor and Ward 1982). Within philosophy, these problems are also often described as “threshold” or “triggering” cases (see, e.g. Hedden 2020 and Nefsky 2019).

number of contributions will fall within a critical zone, and the magnitude of each individual's choice if they do.<sup>2</sup> The decision-theoretic response has its share of detractors. Critics argue that the probability of each contribution being decisive is too low for the expected value calculations to work out in favor of contributing. (See, e.g. Jackson 1997; Lomasky and Brennan 2000; Nefsky 2011; Budolfson 2018; Fullhart and Lord manuscript).

In this entire discussion, which spans over four decades, philosophers have assumed (mostly implicitly) a certain decision-theoretic framework, on which the right act is that which is such that, were you to perform it, good outcomes would most likely obtain.<sup>3</sup> In decision theory, this way of thinking is associated with *causal decision theory* ("CDT"). However, causal decision theory isn't the only option. According to its main rival, *evidential decision theory* ("EDT"), the right act is that which provides the best evidence that good outcomes will obtain.<sup>4</sup>

These theories diverge in certain cases. An act can provide the best evidence for a good outcome by being *correlated* with that outcome, even when the outcome is counterfactually independent of the act in question. Proponents of each theory have offered various cases in which their favored view seems to give the better answer, and the debate is far from settled.<sup>5</sup> Give that

---

<sup>2</sup> Even if, for example, your vote is highly unlikely to be pivotal to the outcome of the U.S. presidential election, it will have major consequences for billions of people if it is.

<sup>3</sup> Authors frequently talk about whether an individual action can "make" or "trigger" a change in the outcome, and about whether individuals' actions are "decisive" or "efficacious." Several theorists have departed from this near consensus in suggesting a *causal involvement* view. On this view, you should perform the act that's most likely to be causally involved in good outcomes (e.g. voting for the winning candidate), even if doing so requires considerable effort and the best outcome is guaranteed to happen anyway (see Goldman 1999; Tuck 2008). This implication strikes me as a decisive reason to reject the causal involvement view, and it hasn't had much uptake (though see Nefsky (2017) for an attempt to develop a version of the view that avoids this implication).

<sup>4</sup> Chignell (2020) is the only author in this literature to discuss evidential decision theory. He argues that individual's actions are correlated in some collective action problems, and that this provides a way of giving an evidentialist justification for individuals making certain sacrifices. However, Chignell only devotes a few pages to this point. His primary concern is to argue that, if one takes oneself to have good reasons not to purchase factory farmed meat, whether or not one's purchase makes a difference to the suffering of animals, then one can be justified in having *faith* in various propositions (including the truth of evidential decision theory) to help sustain one's resolve. In decision theory, Ahmed (2014, 4.6) briefly argues that evidential decision theory can rationalize voting in a national election. In political science, Grafstein (1991) argues that this theory offers a plausible *descriptive* account of voter turnout.

<sup>5</sup> For defenses of causal decision theory, see, e.g. Stalnaker 1981, Joyce 1999, Lewis 1981a, Gibbard and Harper 1978, Skyrms 1980, and Gallow 2021. Defenses of evidential decision theory include Jeffrey 1965, Horgan 1981, Ahmed

it's unsettled, we should consider the prospects of an evidentialist version of the decision-theoretic response to collective action problems.

Drawing upon work by Lewis (1979), I'm going to argue that EDT changes the consequentialist calculus in such cases. When individuals are similar to one another, their actions are non-causally correlated. So even if I don't expect my action to make a difference, I can still expect the outcome to be different, given that I act, because what I do is correlated with what others do and, together, our actions make a difference.

In section 2, I'll give formal characterizations of evidentialist and causal versions of act consequentialism. In section 3, I'll use Newcomb's Problem (Nozick 1969) to illustrate how these views diverge and to discuss some of their underlying philosophical motivations. I'll then show, in section 4, how we can get precisely the same divergence in the Prisoner's Dilemma. Section 5 will extend this argument to show that EDT and CDT give different answers about when individuals should contribute to the provision of lumpy goods: EDT can require individuals to contribute even when CDT does not. I'll focus on the case of voting in a large national election, since many theorists consider the prospects of a consequentialist justification for voting to be hopeless. In section 6, I'll consider several objections to my argument. I'll conclude in section 7 with some tentative reflections on why EDT offers an appealing mix of idealism and pragmatism.

## 2. Causal and Evidential Decision Theoretic Consequentialism

Consequentialism says that an act is right or permissible just in case the value of its consequences is at least as good as that of the consequences of any alternative. The value that determines what

---

2014, and Hare and Hedden 2016. Elga (2020) argues that each theory is intuitively compelling in some cases, and that settling the debate will require one side to provide an error theory of the intuitions that favor the other side's view.

an agent ought to do is a value that is exemplified by whole possible worlds. Each of these worlds can be thought of as one of the possible “total consequences” that the relevant acts can have. To represent consequentialist theories formally, then, we need to introduce a value function  $V(*)$  that measures the degree to which each world  $W$  exemplifies the relevant value.

We can remain fairly neutral about what this value is. Our primary concern is with how to connect worlds to acts in determining each act’s deontic status. We do this with a probability function  $P(*)$ . For an evidentialist-inspired version of act consequentialism, an act’s expected value is a function of each world’s probability of obtaining, conditional on the act’s performance ( $P(W|A)$ ). For an act  $A$  to be right is for it to be the case that for every available alternative act  $B$ , and every possible world  $W$ :

$$\sum_w V(W)P(W|A) \geq \sum_w V(W)P(W|B)$$

For a causalist-inspired version of act consequentialism, we can follow Stalnaker’s formulation of CDT (1981) and say that an act’s expected value is a function of probabilities of the form *if A were performed, W would obtain* ( $P(A \square \rightarrow W)$ ). For  $A$  to be right is for it to be the case that:

$$\sum_w V(W)P(A \square \rightarrow W) \geq \sum_w V(W)P(B \square \rightarrow W)$$

A crucial interpretive question here is how to understand the probability function. How, in other words, do we get the relevant conditional probabilities for EDT, or the probabilities of the relevant counterfactuals for CDT? For our purposes, we can distinguish between three different interpretations.  $P(*)$  can correspond to the *objective* probabilities, the probabilities given the agent’s *available evidence*, or the agent’s *credences* in the relevant propositions. We can then distinguish three senses in which an act can be right according to consequentialism:

- An act is objectively right just in case it maximizes expected value relative to the objective probabilities.
- An act is right in the evidence-relative sense just in case it maximizes expected value relative to the agent's available evidence.
- An act is subjectively right just in case it maximizes expected value relative to the agent's credences.

To remain as ecumenical as possible, I won't privilege one of these notions of rightness over the others in what follows. Rather, I'll discuss my conclusions in light of each of these ways of thinking about the deontic status of an action.

### 3. Newcomb's Problem

A being offers you a choice between taking only an opaque box ("one-boxing") or taking both the opaque box and a transparent box ("two-boxing"). You can see that the transparent box contains one thousand dollars. And you know that the opaque box contains either one million dollars or nothing. Whether it contains a million dollars depends on a prediction that the being made before you make your choice. If the being predicted that you would one-box, it placed the million in the opaque box. But if it predicted that you would two-box, it left the opaque box empty. Suppose that you've seen hundreds of other people play the game with this being, and that its predictions have always been correct. What should you do?

Here's the CDT way of thinking about this case. You know that the million dollars is already in the opaque box or it's not. Nothing that you do will change this. So you should decide whether to one- or two-box based on what would happen, were you to perform each act in each of these possible scenarios. Were the money to be in the opaque box, then you'd be better off taking

the thousand as well. Were the opaque box to be empty, then you'd still be better off taking the thousand. So you should two-box.

EDT sees things differently. Given that you one-box, you'll most likely get a million dollars, because one-boxing is correlated with the being predicting that you'll one-box and placing a million dollars in the opaque box. In contrast, conditional on two-boxing, you'll most likely get only a thousand dollars, since this choice is correlated with a two-boxing prediction. Since you can expect to do much better, given that you one-box, that is what you should do.

The idea underlying EDT is that a rational agent performs whichever act she would most want to *learn* that she had performed, given *all* of her evidence at the time of acting, including the act itself. EDT treats as open anything whose likelihood varies conditional on which act she performs. It asks us to treat an agent as choosing between the *epistemic* upshots of her available acts, rather than their expected *counterfactual* upshots. Someone in Newcomb's Problem, then, is thought of as choosing not only which box to take, but also which prediction was likely made about her own choice.

Evidentialists often motivate their view by observing that followers of EDT end up with more wealth in Newcomb's Problem than followers of CDT (Ahmed 2014, 194; Hare and Hedden 2016). This is known as the "why ain'cha rich?" argument for one-boxing. Defenders of CDT reply that their followers simply have worse opportunities, which they make the best of by taking the guaranteed thousand (see, e.g. Lewis 1981b, 377). We needn't settle this debate. However, we'll soon see that philosophers level a "why ain'cha rich?" charge against act consequentialist reasoning in collective action problems, and that EDT provides act consequentialists with a response.

#### 4. Lewis on the Prisoner's Dilemma

To set the stage for my main argument, let's now consider how CDT and EDT can also come apart in the Prisoner's Dilemma. In this case, You and I each have the option to rat the other one out or to stay silent. We act in complete isolation from one another, but otherwise each of us is aware of all of the relevant facts (so that what each of us objectively ought to do corresponds to what each ought to do subjectively and in the evidence-relative sense). Let's assume a version of *agent-relative* consequentialism that ranks worlds relative to how much money each agent receives in that world. So one world is better than another relative to one agent if that agent has more money in that world, even though that same world is worse relative to another agent who has less money in that world.

To make the parallel between this case and Newcomb's Problem as clear as possible, let's follow Lewis (1979) in giving the prisoners payoffs that correspond to the payoffs in Newcomb's Problem (instead of the payoffs normally used for the Prisoner's Dilemma, which consist (unsurprisingly) in amounts of jail time rather than money). Here, in table form, are our options and our payoffs:

	Me	
	I rat	I don't rat
You	You rat	You get \$1,000 I get \$1,000
	You don't rat	You get \$0 I get \$1,001,000
		You get \$1,001,000 I get \$0
		You get \$1,000,000 I get \$1,000,000

What should each of us do, according to CDT and EDT?

CDT says that each of us should rat. Because You and I act in isolation from one another, our acts are counterfactually independent. When we compare the expected value of me ratting vs. not ratting, we have to hold fixed what you do. First, suppose that you rat me out. Then, if I were

to rat you out, I would get \$1,000, whereas if I were to not rat, I would get nothing. Now suppose that you don't rat me out. Were I to rat, I would get \$1,001,000, whereas not ratting would get me \$1,000,000. Either way, I expect to get an additional \$1,000 by ratting, so I should rat (similar reasoning shows that you should do likewise).

As Lewis (1979) points out, EDT looks at the situation differently. The crucial issue is what I can expect to happen, given that I rat, and what I can expect, given that I do not. Acting in isolation breaks any counterfactual connections between our actions, but they may still be highly correlated. Perhaps you and I are similar in how we approach decision-theoretic puzzles, and we know this about ourselves because we're aware of our respective track records. My choice to rat you out then provides excellent evidence that you will also rat me out, in which case I'll only get \$1,000. However, my choice to not rat is strong evidence that you'll do the same, and I'll receive \$1,000,000.<sup>6</sup> So I should not rat, according to EDT (the same goes for you).

The degree of correlation between our actions needn't be particularly strong for EDT's verdict to differ from CDT's. In this case, not ratting the other person out has a higher expected value (in dollars) for EDT as long as there is more than a .5005 chance that you and I make the same choice.<sup>7</sup> In other words, it could be almost as likely that we'll choose differently, and EDT will still say that we shouldn't rat each other out.

As with Newcomb's Problem, followers of EDT will be rewarded with one million dollars each in the Prisoner's Dilemma, whereas followers of CDT will receive only one thousand dollars a piece. Again, evidentialists see this disparity as the basis for another "why ain'cha rich?"

---

<sup>6</sup> Here, your choice is analogous to the being's prediction in Newcomb's problem. In both cases, I cannot determine the matter in question, but certain possibilities are more likely, given that I act one way rather than another.

<sup>7</sup> To arrive at this number, we express the expected value of not ratting as  $1,000,000 * c + 0 * (1 - c)$ , and the expected value of ratting as  $1,000 * c + 1,001,000 * (1 - c)$ . Here, "c" stands for the degree of correlation between our actions. To know how strong the correlation must be, we solve for the value of c such that  $1,000,000 * c + 0 * (1 - c) > 1,000 * c + 1,001,000 * (1 - c)$ , which is  $c > .5005$ .



argument against CDT. Followers of CDT earn themselves one thousand dollars, when they could have each received one million by not ratting, as EDT prescribes. This disparity, say evidentialists, indicates that EDT is the more plausible theory.

## 5. Lumpy Goods

In the literature on collective action problems, philosophers often echo the “why ain’cha rich?” argument by contending that followers of act consequentialism will inevitably bring about bad, avoidable outcomes, and that the theory is implausible for this reason. For example, Nefsky (2019) writes that there must be something with the argument that ““things will be just as bad whether or not I act in this way, so there's no point in doing otherwise’...[because] when enough people reason in this way, serious avoidable harm results from our voluntary actions” (2).<sup>8</sup> In this section, we’ll see that individuals can avoid causing serious harm while still following act consequentialism, provided that they act based on evidentialist considerations, rather than causalist considerations.

Since I want my discussion to apply to as broad of a range of consequentialist theories as possible, my focus is on collective action problems that arise when goods are provided in a lumpy way. This kind of problem can arise for both agent-relative and agent-neutral consequentialist theories.<sup>9</sup> That is because it stems from the difficulty of coordinating our actions so that each contribution is critical, in the sense that were any given contribution not to have been made, then the good in question would not have been realized (or it would not have been realized to the same extent). The only axiological assumption that we need to make is that a given contribution is not

---

<sup>8</sup> Similarly, Soon (2021) says that “[c]ollective harm is a predictable result of [act consequentialist] reasoning and could be avoided if enough people decided to cooperate” (3354).

<sup>9</sup> In the Prisoner’s Dilemma, followers of CDT end up with only one thousand dollars in part because of their counterfactual thinking, but also because they follow an agent-relative theory that tells each of them to acquire as much money for himself as possible. An agent-neutral version of CDT that tells You and Me to do whatever would yield the highest monetary payoff—regardless of whether it is for You or for Me—requires each of us not to rat.

worth the effort, even in agent-neutral terms, unless it is critical to the realization of the relevant lumpy good.

To illustrate how CDT and EDT can come apart in such cases, I'm going to focus on the decision of whether and how to vote in a national election. Voting is widely considered to be an especially difficult type of action to justify on act consequentialist grounds.<sup>10</sup> It's also a useful test case because we have considerably more evidence about the comparative counterfactual and evidential significance of a single vote than we have regarding, say, the significance of a single purchase of meat.<sup>11</sup> Although voting is an especially useful case to consider for these reasons, the broader lessons of my argument can be applied to other collective action problems.

Now, without further ado, let's consider the following scenario:

**Election:** Daisy is a much better candidate than Donald. Enough people voting for Daisy would result in a reduction in poverty, manageable inflation, reduced risk of war, etc. You are a typical voter, whose decision about whether and how to vote is unlikely to influence anyone else's decision. In order to vote, you will have to miss a couple hours of work driving to and from the polling location.

Are you morally obligated to vote for Daisy? Let's first consider how CDT requires us to think about this situation.

Objectively, you ought to vote for Daisy according to CDT just in case what would happen, were you to vote for her, is better than what would happen if you were to abstain (or vote for Donald). In the vast majority of scenarios, this won't be the case, since voting will be somewhat inconvenient and the outcome will be the same regardless of whether and how you vote. For your

---

<sup>10</sup> The other case that is frequently regarded as especially difficult is the case of climate change. The main points that I make about voting almost certainly apply to this problem as well. However, climate change is impossible to discuss adequately without getting into difficult questions in axiology, which is not my focus here. For a recent discussion of some of the axiological issues raised by problems like climate change, from an act consequentialist perspective, see Hedden (2020, sec. III-IV).

<sup>11</sup> Although the ethics of consumer choice literature is beginning to pay more attention to relevant empirical evidence. See especially McMullen and Halteman 2018.

vote to make a difference, the votes must be evenly split or Donald must only be ahead by one vote (in which case your vote would result in a tie, which might be resolved in Daisy's favor).<sup>12</sup>

The usual consequentialist response here is that you nevertheless ought to vote for Daisy in either the subjective or the evidence-relative sense. Defenders of this response insist that a typical voter isn't completely certain that Donald or Daisy will win regardless of what that voter does, nor does the evidence completely rule this possibility out. Moreover, given how much better things would go if Daisy were to win, these philosophers think that the expected value calculation works out in favor of voting for Daisy (see Barnett 2020).

A serious challenge for this response is that in many elections, one candidate is highly favored going into the election or—more importantly in the U.S., given the electoral college system—one candidate or the other is highly favored in most states. Imagine that you're registered to vote in one such "safe" state, New Jersey, where *The Economist* has stated, based on extensive polling data, that Daisy is going to win between 56 and 64 percent of the vote, and that she has well over a 99 percent chance of winning.<sup>13</sup> Given these numbers, it's overwhelmingly likely that Daisy would win New Jersey, if you were to vote for her, and that she would win, if you were to refrain or to vote for Donald. So whatever happens in New Jersey—and, consequently, whatever happens in the national election—is overwhelmingly likely to be counterfactually independent of what you do. So you should abstain, given the inconvenience of voting, the risk of being involved in a collision on your drive, etc.<sup>14</sup>

---

<sup>12</sup> I'm operating under the simplifying assumption that if a candidate really won by a single vote, that would settle the outcome. One worry about this assumption is that, if the votes were so close, the election might instead be decided by the courts after extensive litigation (Somin 2013).

<sup>13</sup> These numbers come from *The Economist's* predictions for Joe Biden winning New Jersey in the 2020 election. See *The Economist Newspaper* 2020. Of course, not all voters are aware of these statistics. I'm assuming, however, that most voters are at least aware of whether the outcome in their respective states is likely to be close or not.

<sup>14</sup> This is the standard line about voting to change the outcome that economists, political scientists, and philosophers take. See, e.g. Brennan 2020. Not everyone accepts this argument. For recent criticism, see Barnett 2020 (Note that Barnett says nothing about voting in a safe state, nor is it clear how his argument could justify voting in that context.).

Now, a given voter in a “swing” state, where the vote is expected to be much closer, is more likely to cast the decisive vote. On the optimistic side, Edlin, Gelman, and Kaplan (2007) estimate that such a voter has a one in ten million chance of deciding the whole election. Perhaps, given the stakes, this is enough of a chance to make voting required by CDT.

What does EDT have to say about voting? For the evidentialist, whether you ought to vote for Daisy depends on whether she is more likely to win, conditional on you voting for her. This depends not only on the likelihood of you casting the decisive vote, but also on the degree to which your decision is correlated with the decisions of other eligible voters. The degree of correlation depends, in turn, on the extent to which others are like you, at least with respect to the kinds of factors that influence whether and how they vote.

It’s plausible that the factors that affect whether you vote for Daisy or abstain will influence many others, too. There’s considerable homogeneity among voters. For instance, in the United States, people who tend to vote Republican are more likely to be white, male, lacking a college degree, and from rural communities, whereas people who tend to vote Democrat are more often non-white, female, college-educated, and from urban areas (Pew Research Center, 2018). Given these and many other similarities, we would expect the decisions of those who are inclined to vote for Daisy to be correlated with your decision. Indeed, political scientists have amassed considerable evidence that a given individual’s party affiliation and voting decisions are highly correlated with those of other members of demographic groups with which he identifies.<sup>15</sup>

Since whether and how you vote is correlated with whether and how members of large demographic groups vote, your decision has diagnostic value even when the outcome of the election is counterfactually independent of what you do. This point is easiest to appreciate when

---

<sup>15</sup> See Achen and Bartels (2017). I focus on correlations between members of demographic groups because they are fairly easy to observe and have been studied extensively.

we consider voters in safe states. Suppose that 60 percent of the votes in your state will go for Daisy. It can still be the case that your decision is correlated with that of enough voters in swing states for it to turn out that Daisy is significantly more likely win, conditional on you voting for her, than conditional on you abstaining. Moreover, even if CDT requires voters in swing states to vote for Daisy, the case for voting in these circumstances will be stronger on evidentialist grounds, given that the diagnostic value of such a vote will include the possibility that the vote is pivotal as well as the correlation between that vote and other votes.

So far, I've focused on whether your decision is *in fact* correlated with the decisions of other voters. If we're interested in whether individuals ought to vote, given their evidence or their credences, the mere fact that these correlations exist isn't enough. However, some voters—such as those who follow political trends closely—will be aware of these correlations. Moreover, even less informed voters will likely have some idea that their votes indicate how others will vote. There is considerable empirical evidence that people think of their choices as having diagnostic value. Individuals who engage in some behavior tend to regard it as significantly more common than those who engage in an alternative behavior (Mullen et al. 1985).

Psychologists have proposed a variety of mechanisms to explain this phenomenon. These mechanisms include: selective exposure to like-minded others, the salience of the individual's own response, and the individual's emotional attachment to her response being widely shared (Koestner et al. 1995; Krueger 1998).<sup>16</sup> Crucially, though, one such mechanism is inductive inference based on one's own response (Koestner et al. 1995, 222-24; Krueger 1998, sec. III; Quattrone and

---

<sup>16</sup> Another mechanism that may be in play here is the following: individuals think (at least inchoately) that the best response is the more likely response. Individuals also usually give whatever response they think is best. So an individual will tend to think that his response is widespread, because his response and his belief that this response is common share a common cause. Namely, the individual's belief that the response in question is best. I'd like to thank [redacted] for suggesting this explanation.

Tversky 1986, 54-55). That is, individuals regard their own choices as evidence for how others will choose. Moreover, individuals' accuracy in assessing the prevalence of a given response in a larger group tends to improve when they are asked how they themselves would respond, indicating that this information has genuine evidential value (Krueger 1998, sec. III).

So from an objective, subjective, or evidence-relative point of view, EDT's case for voting (for the better candidate) is plausible and more secure than CDT's.

## 6. Objections

Let's now consider several objections, each of which points to a reason why EDT's verdicts about collective action problems might not differ from CDT's.

First, you might think that the evidence that your actions give you about what others will do is insignificant in comparison to all of your other evidence. In a national election, we generally have lots of other evidence to go on: historical trends and analysis, polling data, etc. Surely, you might think, one vote is evidentially insignificant, given all of this other information.<sup>17</sup>

This "too much other evidence" objection doesn't apply to typical U.S. presidential elections, where there are a significant number of swing states in which the outcome is uncertain. Given this uncertainty, and the correlation between voters in swing and non-swing states, your vote gives you some new evidence about the outcomes of the swing states, even if the results for the safe states are all but settled by your other evidence.<sup>18</sup>

---

<sup>17</sup> This objection is related to the so-called "tickle defense" of two-boxing on evidentialist grounds. According to this defense, our actions have little evidential value, because we have introspective access to the sources of our actions (e.g. our beliefs and desires). See Ahmed (2014, sec. 4.3) for critical discussion.

<sup>18</sup> Moreover, many voters are unaware of most of this evidence, so it has no bearing on what they subjectively ought to do.

The objection also fails for a more general reason. It assumes that being well-informed about what others say they're going to do or about what they've done on past, similar occasions diminishes the diagnostic value of your action. It *can* have that effect, of course, by indicating that others are likely to act one way, even conditional on you acting some other way. However, possessing information of this kind may have the opposite effect. It may reveal correlations between your actions and the actions of others. Perhaps you find that, in every past election, you have voted for the candidate who ultimately won. Or maybe polling data indicates that other voters are similar to you in the factors they are considering in their decisions about whether and how to vote.

A second objection for us to consider goes as follows: even if my action gives me some good news that others will do as I do, it could also give me bad news that others will act in a way that undermines these efforts. Perhaps my vote for Daisy is evidence both that like-minded voters are motivated to vote for her by the factors that motivated me, and that these same factors are driving other voters to cast their ballots for Donald. For example, Daisy's party affiliation could motivate me and other party members to vote for her, but lead voters from the opposing party to vote against her. Or her plans to introduce more redistributive taxation may appeal to some voters and turn off others ("that's socialism!"). Maybe, conditional on me voting for Daisy, it's probable that an *equal* number of additional Daisy and Donald supporters will turn out for their respective candidates. In this case, the only good news that I would get by voting would be the possibility that my vote is decisive to the outcome. I'd then have no more reason to vote by EDT's lights than by CDT's.

This kind of scenario is possible. However, it's unlikely that in most cases, my decision to do my part will be correlated with like-minded others doing their part, and correlated with an

equally large number of deviants deciding to act in some way that fully cancels out this good news. First, there may be far more people who are inclined to contribute to the provision of the relevant lumpy good than to undermine others' contributions. Not every collection action problem involves large, competing coalitions of people who are highly polarized. Even in cases such as Election, however, it's hard to imagine that fans of Daisy and fans of Donald are so diametrically opposed that whatever draws Daisy supporters to the polls will bring in exactly the same number of Donald supporters.<sup>19</sup>

It's worth noting that, in some rare cases, EDT will imply that an individual should vote for the worse candidate, since her choice to do so will be correlated with others deciding to vote for the better candidate. Suppose that you're an "undecided" voter who votes more or less at random, without any discernible partisan pattern. Yet you find that you always vote for the losing candidate. Given your track record, the better candidate is more likely to win, conditional on you voting for the worse candidate.

This implication is odd, but it should be unsurprising. If what matters is choosing the action that's the best evidence for the best consequences, then of course it could turn out that an action, such as voting for Donald, might be required because it is more diagnostic of a Daisy victory than a vote for Daisy herself.<sup>20</sup>

---

<sup>19</sup> Certain processes, such as partisan sorting and polarization, may increase the degree to which the level of turnout for one major candidate correlates with turnout for the other. Sorting occurs when individuals move to areas where their fellow partisans live, adopt consumption habits and hobbies that match those of others on their side, and form beliefs on policy that fit what they perceive to be the acceptable set of beliefs. Polarization is the process by which individuals' views become more extreme and are held more confidently (See Talisse 2019, part II for helpful discussion of these processes and their significance for democratic theory.). We would expect sorting to increase both the degree to which the choices of individuals within the same party are correlated, and the extent to which their choices are correlated with different choices from individuals within the rival party. It's hard to say whether the overall effect would be an increase or a decrease in the positive news value of a given vote. Increased polarization likely has both effects, as well.

<sup>20</sup> Political scientists worry that majoritarian democracy is unresponsive to the priorities of persistent minorities, "groups whose members' preferences are positively correlated—and negatively correlated with those of a persistent majority—over a broad range of issues over time" (Abizadeh 2021, 748). Does EDT have any implications for the democratic participation of persistent minorities? This is a huge and difficult topic, especially since the category of



I should stress that EDT won't tell *most* of us to act in this way. It can't turn out that most people's actions are anti-correlated with most other people's actions. If there are ten of us, and my choice between A and B is anti-correlated with everyone else's choice, then for each of the other nine people, his or her choice will be correlated with eight people's and anti-correlated with one person's. So even if EDT directs some people to vote for the bad candidate based on the value of the good candidate getting elected, it won't direct most of us to act in this way.

## 7. Conclusion

To recap, just about everyone in the literature who has engaged with act consequentialist approaches to collective action problems has assumed a causalist view of expected consequences. Many of these theorists think that a typical individual cannot expect to bring about positive change with respect to the relevant outcomes, and are for this reason skeptical of act consequentialist approaches to collective action problems whole cloth. However, the literature has neglected an alternative, evidentialist conception of expected consequences.

The ethics of collective action problems looks different when viewed through the lens of EDT. You can expect that things will turn out better, given that you do your part, because your

---

persistent minority gets applied to many different groups in different political systems. However, let me venture a few suggestions. First, there may be no candidate with any reasonable chance of winning who would advance the interests of members of some persistent minority. In that case, evidentialist considerations won't provide much of a rationale for voting. Whatever you, a member of the persistent minority, do, your decision won't have enough diagnostic value to make it at all likely that your candidate wins. Second, even when members of a persistent minority favor a leading candidate, the factors that motivate their support may be unique to them. If the race is close enough, such that strong support from the persistent minority could decide the election, then EDT may require you to vote for that candidate. If the race won't be decided by the persistent minority, though, then there won't be any point in you voting, since your vote won't have diagnostic value outside of your minority group. Finally, there may be some cases where your choice of whether to vote for your preferred candidate or abstain is correlated with whether members of some larger group vote for their preferred candidate or abstain. Perhaps your decision to abstain from voting for Daisy is evidence both that some members of your group will abstain rather than vote for Daisy, but also that a larger number of Donald supporters will abstain rather than vote for him. In that case, EDT says that you ought to abstain. I'd like to thank an anonymous referee for inviting me to consider EDT's implications for persistent minorities.

action is correlated with enough others doing their part. I illustrated this point by presenting an evidentialist case for voting in a large national election.

You may find it interesting that EDT and CDT come apart in collective action problems, but not see any reason to prefer the former over the latter theory. However, one consideration in EDT's favor—which these cases bring out—is that it stakes out a middle ground between two extremes that we could adopt in thinking about the ethics of collective action problems. At one extreme, we could ask what we should *all* do, and then say that a given individual should simply do her part in this collective response. Call this the “idealistic” approach. Since it completely ignores what others will in fact do, it will often direct individuals to do things that seem pointless or even counterproductive in the actual circumstances. Alternatively, we could follow CDT in focusing exclusively on whether a given individual's contribution would make a difference, given what everyone else will do. Call this the “pragmatic” approach. Because it takes so much as given, there will be many cases in which each agent does her best and yet we fall dramatically short of what we could have achieved collectively.

Unlike the idealistic approach, EDT says that in assessing someone's conduct, we should only consider his options and the epistemic upshots of those options. But in focusing on their epistemic upshots, rather than their counterfactual upshots, EDT's assessment is sensitive to how things would go if others acted like you, at least in circumstances where they are similar to you. You should do your part, provided that enough others will likely do as you do.

## Works Cited

Abizadeh, Arash. 2021. “Counter-Majoritarian Democracy: Persistent Minorities, Federalism, and the Power of Numbers.” *American Political Science Review*, vol. 115, no. 3, 742-56.

- Achen, Christopher and Larry Bartels. 2017. *Democracy for Realists*. Princeton University Press.
- Ahmed, Arif. 2014. *Evidence, Decision and Causality*. Cambridge University Press.
- Barnett, Zach. 2020. "Why you Should Vote to Change the Outcome." *Philosophy and Public Affairs*, vol. 48 no. 4, 422-446.
- Brennan, Jason. 2020. "The Ethics and Rationality of Voting." In Edward N. Zalta (Eds.). *The Stanford Encyclopedia of Philosophy*.  
 <<https://plato.stanford.edu/archives/win2020/entries/voting/>>.
- Budolfson, Mark. 2018. "The Inefficacy Objection to Consequentialism and the Problem with the Expected Consequences Response." *Philosophical Studies*, vol. 176, no. 7, 1711-1724.
- Chignell, Andrew. 2020. "Hope and Despair at the Kantian Chicken Factory: Moral Arguments about Making a Difference." In L. Allais and J. J. Callanan (Eds.). *Kant and Animals*. Oxford University Press. 213-38.
- The Economist Newspaper with Andrew Gelman and Merlin Heidemanns. 2020, November 3. "Forecasting the US 2020 Elections: President." *The Economist*.  
<https://projects.economist.com/us-2020-forecast/president>.
- Edlin, Aaron, Andrew Gelman, and Noah Kaplan. 2007. "Voting as a Rational Choice: Why and How People Vote to Improve the Well-Being of Others." *Rationality and Society*, vol. 19, 219-314.
- Elga, Adam. 2020. "Newcomb University." *Analysis*, vol. 80, no. 2, 212-21.
- Fullhart, Sam and Errol Lord. Manuscript. "Knowledgeable Moral Mathematics."
- Gallow, J. Dmitri. 2021. "Riches and Rationality." *Australasian Journal of Philosophy*, vol. 99, no. 1, 114-129.

- Gibbard, Alan and William Harper. 1978. "Counterfactuals and Two Kinds of Expected Utility." In C. Hooker, J. Leach, and E. McClennan (Eds.). *Foundations and Applications of Decision Theory* (University of Western Ontario Series in Philosophy of Science, 13a). Springer. 125-162.
- Goldman, Alvin. 1999. "Why Citizens Should Vote: A Causal Responsibility Approach." *Social Philosophy and Policy*, vol. 16, 201–217.
- Grafstein, Robert. 1991. "An Evidential Decision Theory of Turnout." *American Journal of Political Science*, vol. 35, no. 4, 989-1010.
- Hare, Caspar and Brian Hedden. 2016. "Self-Reinforcing and Self-Frustrating Decisions." *Nous*, vol. 50, 604–628.
- Hedden, Brian. 2020. "Consequentialism and Collective Action." *Ethics*, vol. 130, no. 4, 530-554.
- Horgan, Terry. 1981. "Counterfactuals and Newcomb's Problem." *The Journal of Philosophy*, vol. 78, no. 6, 331–356.
- Isaacs, Yoav, Adam Lerner, and Jeffrey Russell. Forthcoming. "Counting your Chickens." *Australasian Journal of Philosophy*.
- Jackson, Frank. 1997. "Which Effects?" In J. Dancy (Ed.). *Reading Parfit*. Wiley- Blackwell. 42-53.
- Jeffrey, Richard. 1965. *The Logic of Decision*. University of Chicago Press.
- Joyce, James. 1999. *The Foundations of Causal Decision Theory*. Cambridge University Press.
- Kagan, Shelly. 2011. "Do I Make a Difference?" *Philosophy and Public Affairs*, vol. 39, no. 2, 105-141.

- Koestner, Richard, Gaetan F. Losier, Nicolay M. Worren, Lisa Baker, and Robert J. Vallerand. 1995. "False Consensus Effects for the 1992 Canadian Referendum." *Canadian Journal of Behavioural Science*, vol. 27, 214-25.
- Lewis, David. 1979. "Prisoner's Dilemma is a Newcomb Problem." *Philosophy and Public Affairs*, vol. 8, no. 3, 235-40.
- . 1981a. "Causal Decision Theory." *Australasian Journal of Philosophy*, vol. 59, no. 1, 5-30.
- 1981b. "Why Ain'cha Rich?" *Noûs*, vol. 15, no. 3, 377-80.
- Lomasky, Loren and Geoffrey Brennan. 2000. "Is There a Duty to Vote?" *Social Philosophy and Policy*, 17, 62-82.
- McMullen, Steven and Halteman, Matthew C. 2018. "Against Inefficacy Objections: The Real Economic Impact of Individual Consumer Choices on Animal Agriculture." *Food Ethics*, vol. 1, no. 4: online first.
- Mullen, Brian, Jennifer L. Atkins, Debbie S. Champion, Cecelia Edwards, Dana Hardy and John E. Story. 1985. "The False Consensus Effect: A Meta-Analysis of 115 Hypothesis Tests." *Journal of Experimental Social Psychology*, vol. 21, 262-83.
- Nefsky, Julia. 2011. "Consequentialism and the Problem of Collective Harm: A Reply to Kagan." *Philosophy and Public Affairs*, vol. 39, 364-95.
- . 2017. "How you can Help without Making a Difference." *Philosophical Studies*, vol. 174, 2743-67
- . 2019. "Collective Harm and the Inefficacy Problem." *Philosophy Compass*, vol. 14, no. 4, 1-17.
- Norcross, Alastair. 2004. "Puppies, Pigs, and People: Eating Meat and Marginal Cases." *Philosophical Perspectives*, vol. 18, no. 1, 229-245.

- Nozick, Robert. 1969. "Newcomb's Problem and Two Principles of Choice." In N. Rescher (Ed.). *Essays in Honor of Carl G. Hempel*. Reidel. 114-46.
- Pew Research Center. 2018. "Trends in Party Affiliation Among Demographic Groups." <https://www.pewresearch.org/politics/2018/03/20/1>.
- Quattrone, George and Amos Tversky. 1986. "Self-Deception and the Voter's Illusion." In J. Elster (Ed.). *The Multiple Self*. Cambridge: Cambridge University Press.
- Singer, Peter. 1980. "Utilitarianism and Vegetarianism." *Philosophy and Public Affairs*, vol. 9, no. 4, 325-337.
- Skyrms, Brian. 1980. *Causal Necessity: A Pragmatic Investigation of the Necessity of Laws*. Yale University Press.
- Soon, Valerie. 2021. "An Intrapersonal, Intertemporal Solution to an Interpersonal Problem." *Philosophical Studies*, vol. 178, 3353-70.
- Stalnaker, Robert. 1981. "Letter to David Lewis." In W.L. Harper, R. Stalnaker, and G. Pearce (Eds.). *IFS: Conditionals, Belief, Decision, Chance and Time* (The University of Western Ontario Series in Philosophy of Science). Springer. 151-52.
- Talisso, Robert. *Overdoing Democracy: Why We Must Put Politics in its Place*. Oxford University Press.
- Taylor, Michael and Hugh Ward. 1982. "Chickens, Whales, and Lumpy Goods: Alternative Models of Public Goods Provision." *Political Studies*, vol. 30, 350-70.
- Tuck, Richard. 2008. *Free Riding*. Harvard University Press.