OPTIMIZING POLITICAL INFLUENCE:
A JURY THEOREM WITH DYNAMIC COMPETENCE
AND DEPENDENCE[1]

Thomas Mulligan

Georgetown University

The philosophical question which I wish to address formally in this paper is this: *How should a person influence politics?*

Influence may be exerted in many ways during the deliberative phase (*e.g.* a campaign) preceding judgment aggregation (a vote). For example, prior to an election a voter may write an op-ed to explain a public policy dispute and advance her preferred policy. She may put up signs on her lawn, enticing others to vote as she will. She may donate money to fund a candidate's operations and advertising. And so on.

One natural normative view–which is widespread among the public and, I think, implicitly accepted by many political philosophers–is that so long as a person does her best to get to the truth about politics, she may exert influence without constraint.[2] (Indeed it might be thought that such a voter has a positive duty to exert influence.) On this line, political participation is bad only if it is grounded in selfishness, bias, or ignorance.

[2] Some may hold a stronger view, under which (i) people have a *right* to participate in politics, regardless of their epistemic state; or (ii) unrestrained influence is an essential element of the deliberation which justifies a political process.

I hope to show in this paper that this view is wrong. Even the most altruistic, rational, and informed among us is obliged to limit her political influence.

To demonstrate, I develop a variant of Condorcet's Jury Theorem. We shall consider the perspective of an "opinion leader", who is contemplating how much influence to exert during the deliberation preceding a vote. This influence may affect the process in two ways: (1) It may increase the competence of those whom it reaches, making them more likely to vote correctly; (2) it may increase dependence between voters, reducing the exercise of individual judgment. Because (1) and (2) have opposite effects on the quality of aggregation–the former is good and the latter, bad–the opinion leader faces an optimization problem.[3]

We shall see that it is never optimal for an opinion leader to exert influence without limit. He ought to either refrain from influencing the electorate entirely or exert a finite amount of influence, depending on circumstance. To be sure, when the opinion leader refrains from exerting influence, he reduces the probability that the outcome he believes is correct–which he may have selected in good faith, after careful consideration–will obtain. However, he increases the probability that the political process selects the correct outcome. As all of this is transparent to the opinion leader, he acts wrongly when he fails to constrain his political influence.

In §4 and §5 we consider a number of philosophical ramifications of the model. Among these is that it suggests a resolution to Richard Wollheim's (1962) "paradox in the theory of democracy" (*aka* "Wollheim's paradox"). As we shall see, democracy (like other forms of collective decision-making) does not require that a person hold "two incompatible beliefs about what ought to be done" (Ewin 1967: 356) and, thus, is not a "sadly confused idea".

---

[3]Technically, dependence is bad from the point-of-view of aggregation quality only if voters are *competent*; that is, more likely than a coin flip to vote correctly (see §2). No assumptions about voter competence are imposed in this paper's model.

1 Preliminaries

Before turning to formal matters, let me say a few words about three discussions which touch on the philosophical point of this paper.

There is, first, the enormous literature on the epistemology of disagreement. The core question there is how a person ought, rationally, to modify his belief in a proposition X upon learning that someone else disagrees about X.[4] Two camps have emerged: the "steadfast" camp, which holds that maintaining one's belief is permissible, and the "conciliationist" camp, which demands some degree of belief revision.[5]

My concern, however, is not (solely) what a person ought to *believe* about political matters, but how a person ought to *act* during political deliberation. There are a number of interesting issues at the intersection of the epistemology of disagreement and collective decision-making which have not yet been fully explored. For example: (i) Even if conciliationism is generally correct, should a person conciliate when participating in collective decision-making? (The problem being that that would increase dependence between participants, thereby reducing the quality of the group decision.)[6] (ii) How should a person modify his belief when he disagrees not with another person, but with the outcome of a judgment aggregation (when his disagreement is, in some sense, with a group)?

The second work that bears on this topic is a series of papers by Timothy

---

[4]For an introduction to the epistemology of disagreement, see Christensen 2009 and Feldman 2007.

[5]On the steadfast position, see, *e.g.*, Bergmann 2009, Kelly 2005, and van Inwagen 2010. On conciliationism, see Christensen 2007, Elga 2007, and Feldman 2007. I offer my (generally conciliatory) thoughts in Mulligan 2021.

[6]Example: There is a three-member civil jury. Two jurors believe, with credence 0.4, that the facts and the law favor the plaintiff. The third juror believes this with credence 0.9. So they "conciliate", and all adopt the average credence of 0.57. Before conciliating, the respondent would prevail, two votes to one; after conciliating, the plaintiff wins unanimously. It is unclear that this is an epistemic improvement, to say nothing of justice. (*N.B.* in American civil cases the standard of proof is a "preponderance of the evidence" and so the two possible outcomes are symmetric. For criminal cases, with the much higher "beyond a reasonable doubt" standard, this is not the case.)

Feddersen and Wolfgang Pesendorfer (*viz.* their 1996, 1997, and 1999).[7] They provide a game-theoretic justification for abstaining from voting. The basic idea is this: Voters only care about casting a vote because that vote might end up "swinging" the election. Thus, analogous to the "winner's curse" in auction theory, one ought to condition one's vote on what must be true about the world if the election is, indeed, on a knife's edge. If it is, according to Feddersen & Pesendorfer's model, then an "uninformed" voter will strictly prefer abstention in order to allow an "informed" voter to cast the deciding vote.

In contrast, we shall examine how a voter attempts to influence others' votes. In the model of this paper, it is assumed that all voters will cast a vote when deliberation concludes. Furthermore, in Feddersen & Pesendorfer's model, when a voter abstains it is because it is rational *by his own lights.* In the model here, "abstention" (*i.e.* refraining from influencing the political process) does not have that effect. Quite the opposite: It *reduces* the the probability that the voter's preferred outcome will obtain. Thus the justification for restraint cannot be based on self-interest, as it is for Feddersen & Pesendorfer. A normative standard must be invoked.

Third, I wish to point out a remark by Christian List and Kai Spiekermann, who (in discussing Goldman 1999) say that there is an "important ambiguity in existing democratic balloting procedures". Namely: Should voters in the voting booth seek to answer "(1) Which candidate or party would be best *from my own perspective*?", or "(2) Which candidate or party would be best *from the perspective of society as a whole*, allowing for the fact that this perspective may differ from my own?" (2016: 229-30).

Here we consider an analogous ambiguity, between (1) "What amount of political influence should I exert to maximize the probability that the correct candidate (party, policy, *etc.*) *as I see it* is chosen?", and (2) "What amount of

---

[7]See also Austin-Smith 1990.

political influence should I exert to maximize the probability that the *political process* will choose the correct candidate (party, policy, *etc.*)?"

On to the formalism. Our general framework is collective decision-making, the core result of which is Condorcet's Jury Theorem (Condorcet 1785). It has found application in juridical matters (as was its intention), management, and also–indeed most successfully and most frequently–in democratic theory (see, *e.g.*, Landemore 2013).

Common to all these contexts is that there is a group of people who have a common goal but disagree about how to pursue that goal. All 12 jurors may wish to convict a defendant if and only if he is guilty, but they disagree about whether this is the case; some believe he is guilty and others, not. Or, we have a group of voters who all want to choose the *just* policy, but who disagree about which of the options facing them is in fact just. The question, then, is how to optimally aggregate these differing opinions to maximize the probability that the shared goal obtains.

We consider an *electorate* of n *voters*, where $n = 2m + 1$. Observe that n is odd. (I shall use political terminology since I am interested in political matters; of course the model holds across circumstances of collective decision-making.) Voters face a dichotomous choice–they must, say, cast a vote for (exactly one of) policy A or B. Associated with each voter $v_i \in \{v_1, v_2, ..., v_{2m}, v_n\}$ is a *competence* $c \in [0, 1]$, which is that voter's probability of casting a correct vote. We assume that this value is shared; no voter is more or less competent than any other. We also assume that the electorate will aggregate votes using simple majority rule–"one person, one vote". For now we assume that votes are statistically independent.

Condorcet showed that, given these assumptions, the probability that the

5

electorate will choose correctly is:

$$U = \sum_{i=m+1}^{n} \frac{n!}{i!(n-i)!} c^i (1-c)^{n-i}. \tag{1}$$

So long as $c > 0.5$–that is, voters are better than a coin flip at choosing correctly–equation (1) evinces two key features. First, U is strictly increasing in n. From the perspective of good group decision-making, the more voters you add, the better. Second, $U \to 1$ as $n \to \infty$.

It is not hard to see that some of the assumptions of equation (1) are dubious in practice. Most problematic are the assumptions of homogeneous competence and statistical independence.[8]

The assumption of statistical independence is more troublesome. I say this for two reasons. First, the case of heterogeneous competence has been well dealt with, theoretically. Bernard Grofman *et al.* (1983) and Shmuel Nitzan & Jacob Paroush (1982) demonstrated that, in such cases, a uniquely optimal decision rule, not necessarily simple majority rule, is yielded by any set of voter competences.[9]

Second, it is my opinion (Mulligan 2018a) that human beings are not, in many cases, so different in our abilities. In these cases, the assumption of homogeneous competence works (even if homogenous competence is not strictly speaking true, simple majority rule may remain optimal).

But statistical independence is a big problem. I can think of no real-world case in which it holds, even roughly. Jurors receive common information during

---

[8]This is not to say that the other assumptions may never be violated in practice. But they are often complied with. Many real-world decisions are naturally dichotomous, or they are not naturally dichotomous but are broken down into dichotomous steps for practical reasons. And simple majority rule is ubiquitous (I offer an alternative in Mulligan 2018b). In any case, work has been done extending Condorcet's theorem to deal with these limitations; on the relaxation of dichotomous choice, *e.g.*, see Hummel 2010, Lam & Suen 1996, List & Goodin 2001, Miller 1996, and Paroush 1990.

[9]In particular, the optimal decision rule may be identified by weighting each vote in proportion to the log-odds of that voter's competence. Other relevant literature includes Grofman & Feld 1983, Nitzan & Paroush 1984 and 1985, and Paroush 1998.

a trial. The directors of a company review the same financial documents before deciding whether to execute a merger. Voters do what concerns us here: try to influence each other.

Indeed, the assumption of statistical independence is discordant with the collective decision-making framework. The whole point is that we want to harness the "wisdom of the crowd" in support of a shared goal. But we can only have a shared goal in the first place if we have something in common which renders that goal desirable to us all. Talk about, for example, a statistically independent jury is absurd, because we constitute a jury to evaluate the facts of some *particular* (putative) crime. If 12 jurors evaluate 12 different sets of evidence, then they are not a jury at all. In this way, perhaps independence may be ruled out *a priori*.[10]

2 The model

We are interested in a "micro" question as well as a "macro" question. The former is, "How does an individual voter decide how much influence to exert?" The latter is, "How well does the electorate do under the various choices that voter might make?"

Our starting point is a generalization of classical Condorcet which is due to Philip Boland (1989) and Boland *et al.* (1989).[11] Let $X_i \in \{X_1, X_2, ..., X_{2m}, X_n\}$ be a random variable, equal to 1 just in case v_i votes correctly and 0 otherwise.

We incorporate dependence by supposing that $X_1, ..., X_{2m}$ are correlated with X_n, but that $X_1, ..., X_{2m}$ are independent when conditioned upon X_n.

---

[10]One might think that this problem can be bypassed by shifting from a brute independence assumption to independence *conditional upon common information* (like evidence). Unfortunately, such a shift imperils other necessary assumptions (like the minimal competence assumption)–see Dietrich & Spiekermann 2013a, 2013b, and 2020 and Ladha 1993.

[11]For other models and discussions of dependent voting, see Berend & Sapir 2005 and 2007; Berg 1993 and 1996; Dietrich & List 2004; Kaniovski 2010; Kaniovski & Zaigraev 2011; Ladha 1992, 1993, and 1995; List & Pettit 2004; Nitzan & Paroush 1985; Peleg & Zamir 2012; Shapley & Grofman 1984; and Zaigraev & Kaniovski 2013.

That is, v_n is the sole source of dependence within the electorate. He is our *opinion leader*.

Let the dependence $\rho := \text{corr}(X_i, X_n)$, where $i \neq n$. We postulate the conditional probability $Pr(X_i = 1|X_n = 1) = c + \rho(1 - c)$. That is, v_i votes correctly if he either (i) decides correctly on his own or (ii) decides incorrectly on his own but is induced to follow the (correct) decision of v_n. Similarly, $Pr(X_i = 1|X_n = 0) = c - \rho c$. We obtain as a consequence $\text{corr}(X_i, X_j) = \rho^2$ when $i \neq j$ (for $i, j \neq n$).

The probability that this electorate will choose correctly is:

$$U(c, \rho) = c\beta_{m|2m}(c + \rho(1 - c)) + (1 - c)\beta_{m+1|2m}(c - \rho c), \qquad (2)$$

where β is (a variant of) the incomplete beta function

$$\beta_{k|n}(\lambda) = \frac{n!}{(k - 1)!(n - k)!} \int_0^\lambda x^{k-1}(1 - x)^{n-k}dx. \qquad (3)$$

The function β provides the probability that at least k out of n votes will be correct, given that each vote is correct with probability λ. Thus, equations (2) and (3) tell us that the electorate chooses correctly if either (i) the opinion leader votes correctly and at least half of the remaining $2m$ voters vote correctly, or (ii) the opinion leader votes wrongly but a strict majority of the remaining voters (and thereby a majority of voters *simpliciter*) votes correctly.

Boland *et al.* prove the intuitive results that U is (i) decreasing in ρ for $c > 0.5$, (ii) increasing in ρ for $c < 0.5$, and (iii) equal to 0.5 for all ρ when $c = 0.5$. Monotonicity in c also holds in this generalization of classical Condorcet:

Theorem 1. *U is strictly increasing in c for c, $\rho \in$ [0,1].*

Proof. From equations (2) and (3) we have:

$$\frac{\partial U}{\partial c} = \beta_{m|2m}(c + \rho(1-c)) - \beta_{m+1|2m}(c - \rho c)$$
$$+ \frac{(2m)!}{(m-1)!m!} \times (c - \rho c)(c + \rho(1-c))^{m-1}(1 - (c + \rho(1-c)))^m$$
$$+ \frac{(2m)!}{(m-1)!m!} \times (1 - (c + \rho(1-c)))(c - \rho c)^m(1 - c + \rho c)^{m-1}.$$

U is defined and continuous on [0,1], so we may restrict ourselves to $c \in$ (0,1).

Observe that $(c + \rho(1 - c))$ and $(c - \rho c)$ take on values between 0 and 1, and that the former is greater than or equal to the latter because $\rho \geq 0$. Moreover, the integrands of $\beta_{m|2m}$ and $\beta_{m+1|2m}$ are mirror images about 0.5 for $0 \leq x \leq 1$, with the former greater than the latter for $0 < x < 0.5$. It follows that $\beta_{m|2m}(c + \rho(1-c)) - \beta_{m+1|2m}(c - \rho c) > 0$.

It may further be verified that the factorial term, $(1 - (c + \rho(1 - c)))$, and $(1 - c + \rho c)$ are non-negative. $\qquad\square$

Now, in Boland *et al.*'s model, ρ and c are exogenous to v_n. If v_n could select them, he could maximize the probability that the correct choice was collectively made by maximizing c and minimizing ρ. If he wanted his preferred policy to be selected by the electorate–*regardless of whether it was in fact correct*–then he would select $\rho = 1$, assuring that this would be the case.

But this is unrealistic, in politics and in other forms of collective decision-making. An opinion leader cannot affect competence or dependence directly. But he can affect them indirectly, by exerting *influence*. For example, he can write an op-ed, which, in turn, could affect either the electorate's competence, or the dependence within it, or neither, or both. The same is true for other forms of political influence (campaigning, giving speeches, advertising, *etc.*) This is the dynamic we wish to model.

Let $i \in [0, \infty)$ be the amount of influence v_n exerts. Let $c_{init} \in [0, 1]$ be the competence of voters *before* v_n exerts influence; that is, c_{init} represents competence prior to opinion leader effects. Then the updated competence of voters, c–that is, their competence *after* the opinion leader affects them–may be given as the following logistic function of i:

$$c(i) := \frac{2 - 2c_{init}}{1 + e^{-ai}} + 2c_{init} - 1, \tag{4}$$

where the parameter $a > 0$ may be interpreted as the strength of our opinion leader's influence on competence. The following figures give examples of i's effect on competence for two values of c_{init}:
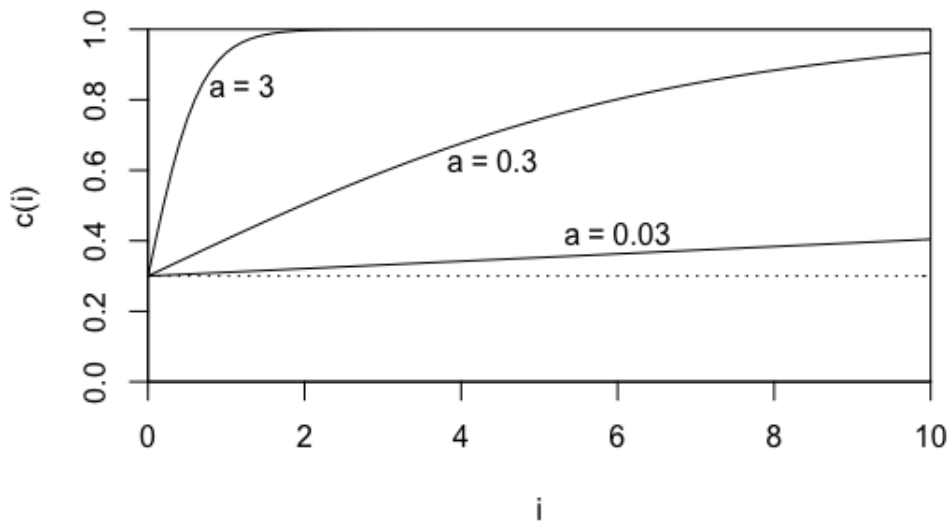

Figure 1: Effect of opinion leader's influence on competence, $c_{init} = 0.3$
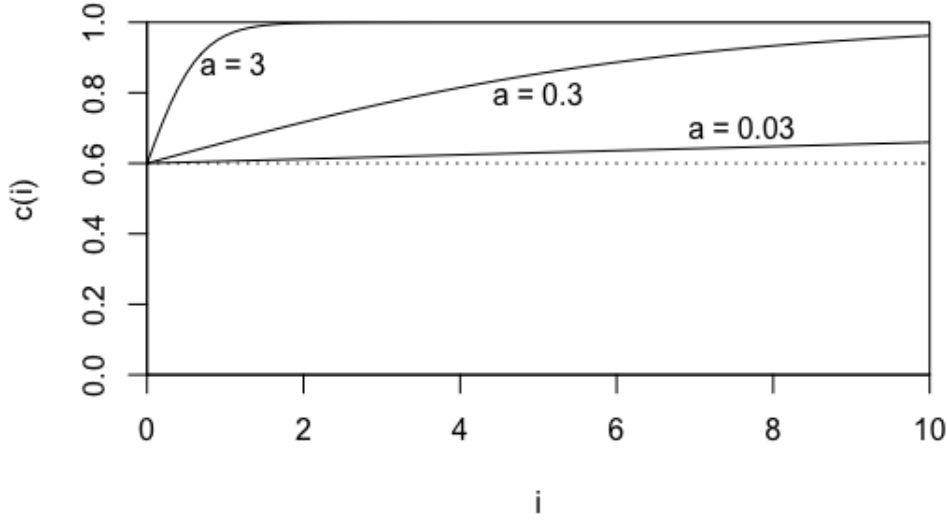
Figure 2: Effect of opinion leader's influence on competence, $c_{init} = 0.6$

The following result follows immediately:

Theorem 2. *U and c are strictly increasing in c_{init}.*

Proof. We have:

$$\frac{\partial}{\partial c_{init}} \left(\frac{2 - 2c_{init}}{1 + e^{-\alpha i}} + 2c_{init} - 1 \right) = -\frac{2}{1 + e^{-ai}} + 2.$$

Because $1 + e^{-ai} \in (1, 2]$, this expression is positive, and because U is a strictly increasing function of c (Theorem 1), by composition U is a strictly increasing function of c_{init}. □

In a similar way we may define the dependence, ρ, as a function of i:

$$\rho(i) := \frac{2}{1 + e^{-bi}} - 1 = \tanh\left(\frac{bi}{2}\right), \tag{5}$$

11

where the parameter $b > 0$ is, again, interpreted as the strength of our opinion leader's influence–but this time his influence on dependence rather than competence.
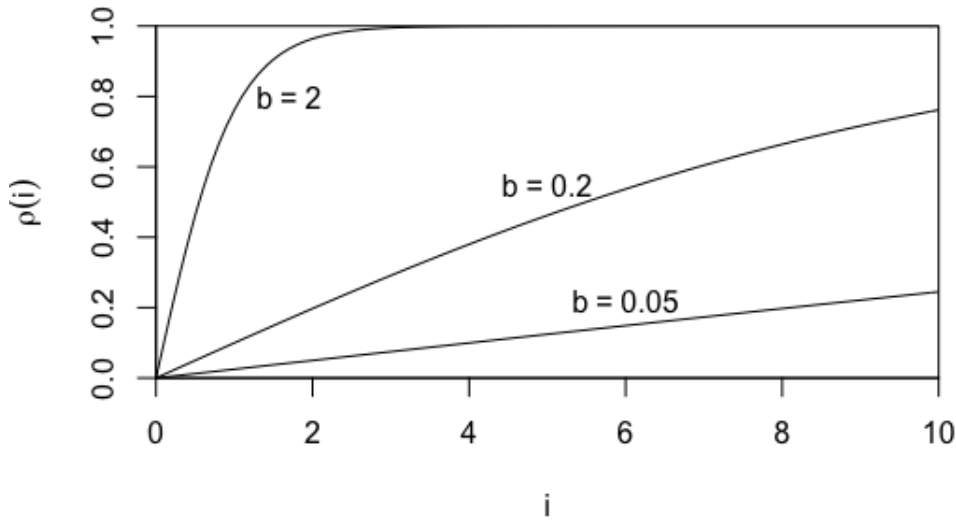

Figure 3: Effect of opinion leader's influence on dependence

Note that i, a, and b have no definite real-world interpretation. If, for example, $3a = b$, we cannot say "the opinion leader's influence creates three times as much dependence as competence". Similarly, an opinion leader who devotes her every waking moment to influencing the political deliberation may be represented by $i = 1$ just as well as $i = 1,000,000$. These specifications can only be made after econometric study. (Simplifying matters somewhat is the fact that in equations (4) and (5) i always appears coupled to a or b. There are two degrees of freedom, not three.)

Now, what can be said is that c and ρ are increasing and concave. This

represents diminishing (but positive) marginal returns to influence. This feature is sensible. It stands to reason that the opinion leader will introduce the most impactful form of influence he can. There could be exceptions. Consider, for example, a demagogue whose influence only produces ρ. It may take great effort for him to build up a base of support, during which time he barely affects the political process. But once a critical level of support is gained, his impact (*i.e.* creation of ρ) increases rapidly. In such a case, the functional specification of equation (5) may not be appropriate.[12]

Before moving on, I recapitulate the model. We imagine ourselves to be in a classic Condorcet-type situation, in which, in particular, voters are independent (and thus uncorrelated). (The situation is actually a bit more general than classic Condorcet, since we make no assumptions about the minimal competence of voters.) We then suppose, as in Boland *et al.*'s approach, that there is a single source of dependence: the opinion leader.

But rather than simply *stipulate* this dependence, we consider a plausible mechanism for its creation–the exertion of influence. We may remain agnostic about the form that this influence takes, and the reasons why it affects dependence and competence within the electorate. Nevertheless, here are two plausible examples: (i) The opinion leader gives a passionate argument, based in his personal experience, for why policy A rather than B should be chosen. His perspective enriches voters' understanding of how the policies affect people, making them more likely to vote correctly. At the same time, some voters are so awed by his rhetoric that they abandon their own judgment to "follow the [opinion] leader" (we may interpret the correlation variable ρ as the probability

---

[12] A natural question is whether the core results of this paper can be obtained with general functions c and ρ, without choosing specific functional forms. General functions would be constrained as follows: Both at least C^1 and strictly increasing in i; $c(i) = c_{init}$ for $i = 0$, $c(i) = 1$ in the limit as $i \to \infty$; $\rho(i) = 0$ for $i = 0$, $\rho(i) = 1$ in the limit as $i \to \infty$. The answer is no. Some functions approximating certain fixed curves, which satisfy these constraints, may yield an infinite number of optima rather than the unique optimum obtained here. I thank Willie Wong for identifying this counterexample.

that a voter does this). (ii) The opinion leader provides the other voters with new evidence, relevant to the problem, and then concludes that the evidence supports A rather than B. The voters benefit, epistemically, from the evidence. But rather than making up their own minds about whether the evidence supports A or B, some voters simply assume that the opinion leader has gotten things right, and so vote for A. (As we will see in §3, the argument of this paper holds even if the opinion leader is epistemically superior to other voters.)

We might also consider dependence using the language of "common causes", introduced by Franz Dietrich & Kai Spiekermann (2013a, 2013b, and 2020) (see also Reichenbach 1956). Observe that two voters, v_i and v_j $(i, j \neq n)$, are correlated despite having no interaction with each other. The correlation arises *via* a common cause, v_n. The opinion leader thus serves as both a direct cause of dependence (between himself and other voters) and a covert, common cause of dependence (between the other voters).

Dietrich & Spiekermann's principal point is that the independence assumption of classical Condorcet, which is obviously no good (§1), cannot be recovered by simply conditionalizing on the common causes of votes (*cf.* n. 11). The reason this is so, roughly, is that the more fine-grained detail we conditionalize on, the less plausible the necessary assumption about voter competence be.

The model of this paper is compatible with Dietrich & Spiekermann's account. Observe that not all the causes of correlation need "run through" the opinion leader. There may be myriad common causes (evidence, arguments, environmental factors) which affect all n voters. These common causes give rise to the baseline competence c_{init} and create dependence. But then we conditionalize on those common causes, eliminating all dependence. To be sure, this may harm our electorate for the reasons Dietrich & Spiekermann give: If these fine-grained common causes are misleading, then c_{init} will be low. This model

is concerned with a subsequent question: Given some baseline competence and no dependence, what happens when a voter can influence his peers?

We are considering one pattern of dependence which arises as a result of voter influence. We can easily imagine other plausible patterns (imagine, *e.g.*, that voters are divided into polarized blocs–say, Democrats and Republicans– and the opinion leader is a member of one of these blocs). I have chosen this approach to dependence for the following reasons: First, it clearly illustrates the philosophical points of this essay. Second, it is a natural way to extend extant work on jury theorems with dependence. Again, rather than just saying how much dependence exists, we are endogenizing dependence and incorporating it into a plausible real-world dynamic. The third reason is pragmatic: The presence of interdependencies renders calculations difficult or even impossible.[13] Indeed, as we shall see, even this simple approach to dependence must be studied numerically, except for special cases.

We may sort opinion leaders into one of four types:

(1) "Anonymous citizen", for whom both a and b are small. No matter how much influence this opinion leader exerts, dependence is not created and there is little new competence generated: $\rho \approx 0$ and $c \approx c_{init}$. Most of us fall into this type, I suppose, owing to lack of resources (monetary, intellectual, *etc.*) or circumstance. When $a \to 0$ and $b \to 0$, U simplifies to the classic Condorcet

---

[13]As one example of the difficulties dependence may introduce, consider a simple case of three voters whose competence and pairwise correlation are known precisely. How likely is it that this electorate will choose correctly? It is impossible to know. These data fail to specify a unique joint distribution (although they do constrain potential distributions). The most common representation of the joint distribution is that given by Bahadur (1961) and Lazarsfeld (1956), which includes $2^n - n - 1$ correlation parameters. The challenge of applying this representation is that some of these parameters (often the higher-order correlations) are unavailable, and generally cannot be ignored (*i.e.* set to 0). To deal with these limitations, Van Der Geest (2005) uses a maximum entropy method, which can be implemented numerically, to infer higher-order correlations. Kaniovski (2008a) obtains the relevant probabilities by identifying the distribution that is "closest" (in a least squares sense) to that in which voters are assumed to be independent. This approach also may be implemented numerically, and Kaniovski obtains an analytic solution for the special case of homogeneous competence. (Numerical examples of this approach are provided in Kaniovski 2008b.)

Jury Theorem. (The same is true when $i = 0$.)

(2) "Disinterested educator", for whom a is large and b is small. Here our opinion leader's influence serves to improve the quality of others' decision-making but does not bind their decisions to his own. Think, for example, of an economist who writes an op-ed explaining deadweight loss during a debate about tax policy. This is the ideal type from the point-of-view of good collective decision-making.

(3) "Fox News", for whom a is small but b is large. Here the opinion leader adds nothing epistemically useful to the deliberation but does induce individuals to abrogate their own judgments to "follow the leader". This is the worst case.

(4) "Knowledgeable partisan", for whom a and b are large. In this case, which is a common one, the opinion leader's influence increases both competence and dependence within the electorate. This opinion leader, especially, faces the non-trivial task of choosing an i that maximizes U.

We now restate U in light of equations (4) and (5).

$$U(i) = c_{init}\beta_{m|2m}(c + \rho(1 - c)) + (1 - c_{init})\beta_{m+1|2m}(c - \rho c), \qquad (6)$$

where

$$
\begin{aligned}
c + \rho(1 - c) &= \frac{2 - 2c_{init}}{1 + e^{-ai}} + 2c_{init} + \frac{2}{1 + e^{-bi}} - 2 \\
&\quad - \frac{(e^{ai} + 2c_{init} - 1)(e^{bi} - 1)}{(e^{ai} + 1)(e^{bi} + 1)} \\
&= \frac{2(e^{ai} + 2c_{init} - 1)}{(e^{ai} + 1)(e^{bi} + 1)} + \tanh\left(\frac{bi}{2}\right), \text{and} \\
c - \rho c &= \frac{2 - 2c_{init}}{1 + e^{-ai}} + 2c_{init} - 1 \\
&\quad - \frac{(e^{ai} + 2c_{init} - 1)(e^{bi} - 1)}{(e^{ai} + 1)(e^{bi} + 1)} \\
&= \frac{2(e^{ai} + 2c_{init} - 1)}{(e^{ai} + 1)(e^{bi} + 1)}.
\end{aligned}
\qquad (7)
$$

This is a single-variable function of i; note in particular that our opinion leader takes the general level of competence, c_{init}, as given.

Now consider an opinion leader who wants more than anything that the electorate select the *just* policy (we shall assume that the just policy is the correct policy). (We can easily imagine other forms of behavior; *e.g.* an opinion leader who desires injustice rather than justice, or who doesn't care about justice at all and simply wants whatever policy will benefit himself to be chosen.) How much influence will he exert?

Again, there is an ambiguity here. The opinion leader may think carefully about justice; educate himself tirelessly about this policy dispute; work to purge all biases from his reasoning; settle on a policy choice–say, A; and then select an i that maximizes the probability that A is chosen. It is not hard to see that in this case the opinion leader will exert i without limit.

Or, the opinion leader can choose an i that maximizes the probability that the electorate chooses the just policy–*whatever it be*. This opinion leader selects:

$$\underset{i \in [0,\infty)}{\operatorname{argmax}} U(i).$$

In both cases, we have an opinion leader who may be fairly characterized as acting out of a concern for, and in accordance with, justice. But we should prefer the latter behavior, for reasons I shall explain.

One might object that it is implausible that the opinion leader would know exactly how various choices of influence would affect competence and dependence within the electorate. He might not know c_{init}, for example. He might be unable to specify c or ρ exactly.

I agree that that is implausible, but it does not affect the core conclusions which this model suggests, and which will be discussed in detail in §§4-5. These conclusions include: (i) When it comes to political influence, there is a trade-

off between its good, competence-enhancing effects and its bad, dependence-producing effects; (ii) individuals should think carefully about *how* to exert political influence, preferring forms that enhance competence and avoiding forms that produce dependence; and (iii) many of our opinion leaders exert too much influence, to the detriment of justice and perhaps their own goals. None of this requires knowledge of how to find, exactly, the optimum.

Dietrich says that "research should concentrate on jury theorems with justifiable premises. A good indicator for whether premises are justified is whether the conclusion is prima facie plausible." (2008: 67). I regard these conclusions as not just *prima facie* plausible but to some extent validated by contemporary political behavior and the performance of actual aggregative systems.

3   Analysis

Figure 4 illustrates the opinion leader's problem for an example set of parameters. The horizontal axis denotes the influence i exerted by the opinion leader and the vertical axis denotes $U(i)$, the probability that the electorate chooses correctly. Note three things. First, if the opinion leader refrains from exerting influence ($i = 0$), then the probability of a correct collective decision is that given by classic Condorcet–in this case, 75%.

Second, if the opinion leader exerts influence without limit, then, eventually, every voter's judgment is perfectly correlated to his own, and the probability that the group chooses correctly is simply the probability that the opinion leader chooses correctly: 60%.

Third, as is clear from the figure, neither extreme is desirable; U reaches an optimum at $i = 0.74$. That is the best balance of competence enhancement (which is *ceteris paribus* good) and dependence creation (*ceteris paribus* bad), and yields a collective decision which has a 90% chance of being correct.
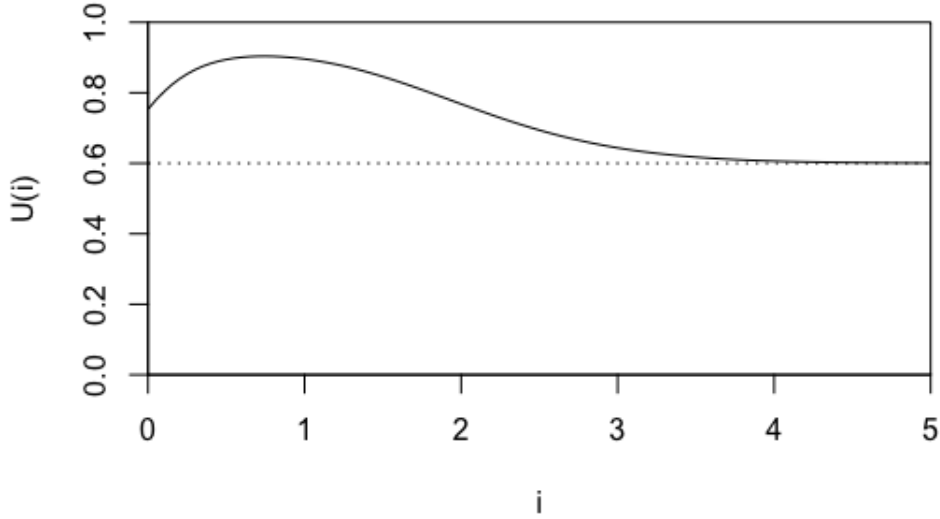
18

Figure 4: Effect of opinion leader's influence on the quality of the collective decision ($c_{init} = 0.6, m = 5, a = 1.5, b = 0.5$)

It is not generally true that there will be an interior optimum, as there is here. For some parameter specifications (*e.g.* "Fox News"), the ill effects of dependence swamp the competence enhancement that i provides. In these cases, the optimal choice for the opinion leader is to exert no influence. On the other hand, as we shall see, it is never optimal to exert influence without limit. And this is plainly true for the typical case, where $c_{init} \in (0.5, 1)$: The opinion leader can always choose $i = 0$, for which $U > c_{init}$ owing to the monotonicity in n feature of classic Condorcet.

To stress, when I say such-and-such influence is the "optimal choice" for an opinion leader, I am referring to an opinion leader who seeks to maximize the probability that the upcoming aggregation selects the just policy. But if the opinion leader's preference is to establish justice *as he sees it*, then it is not

optimal.

Both types of opinion leaders may fairly be said to be concerned with justice. Both might above all else want to create a just society. But the former recognizes the limitations of his own insight into the truth and comports himself during the deliberation accordingly. Unfortunately, much contemporary political activity comports with the latter description (to be discussed in §4).

We may obtain some analytic results. As an intermediate step, we compute:

$$\frac{d}{di}(c + \rho(1-c)) = -\frac{(4c_{init} - 4)((a+b)e^{ai+bi} + ae^{ai} + be^{bi})}{(e^{ai}+1)^2(e^{bi}+1)^2}, \text{and}$$

$$\frac{d}{di}(c - \rho c) = -\frac{2be^{2ai+bi} + (4ac_{init} + 4bc_{init} - 4a)e^{ai+bi} + (4ac_{init} - 4a)e^{ai} + (4bc_{init} - 2b)e^{bi}}{(e^{ai}+1)^2(e^{bi}+1)^2}.$$

Differentiating under the integral sign yields:

$$\frac{dU}{di} = \frac{(2m)!}{(m-1)!m!}\left[c_{init} \times -\frac{(4c_{init}-4)((a+b)e^{ai+bi} + ae^{ai} + be^{bi})}{(e^{ai}+1)^2(e^{bi}+1)^2} \right.$$

$$\times \left(1 + \frac{4c_{init}-4}{(e^{ai}+1)(e^{bi}+1)}\right)^{m-1}$$

$$\times \left(\frac{4-4c_{init}}{(e^{ai}+1)(e^{bi}+1)}\right)^m$$

$$+$$

$$\left(1 - c_{init}\right) \times -\frac{2be^{2ai+bi} + (4ac_{init} + 4bc_{init} - 4a)e^{ai+bi} + (4ac_{init} - 4a)e^{ai} + (4bc_{init} - 2b)e^{bi}}{(e^{ai}+1)^2(e^{bi}+1)^2}$$

$$\times \left(\frac{2e^{ai} + 4c_{init} - 2}{(e^{ai}+1)(e^{bi}+1)}\right)^m$$

$$\left. \times \left(1 - \frac{2e^{ai} + 4c_{init} - 2}{(e^{ai}+1)(e^{bi}+1)}\right)^{m-1} \right].$$

$$(8)$$

We make one assumption: If there are multiple optima, then the opinion leader selects the one with the lowest i. That is, we assume that exerting influence is costly.

We consider the boundary cases:

20

Theorem 3. *When $c_{init} = 1$, the opinion leader chooses $i = 0$. When $c_{init} = 0$, the opinion leader chooses the unique i such that $-2be^{2ai+bi} + 4ae^{ai+bi} + 4ae^{ai} + 2be^{bi} = 0$.*

Proof. The first case is trivial: When $c_{init} = 1$, equation (8) is identically 0, and so is at a maximum under any choice of i. Under our assumption, the opinion leader chooses $i = 0$.

The second case is $c_{init} = 0$. After simplifying equation (8), and because the factorial term is positive, we need only examine the following expression:

$$-\frac{2be^{2ai+bi} - 4ae^{ai+bi} - 4ae^{ai} - 2be^{bi}}{(e^{ai}+1)^2(e^{bi}+1)^2}$$
$$\times \left(\frac{2e^{ai} - 2}{(e^{ai}+1)(e^{bi}+1)}\right)^m$$
$$\times \left(1 - \frac{2e^{ai} - 2}{(e^{ai}+1)(e^{bi}+1)}\right)^{m-1}.$$

The second term equals 0 only when $i = 0$. The third term is positive: Because $\frac{2e^{ai}-2}{(e^{ai}+1)(e^{bi}+1)}$ is decreasing in b, we may consider this term with $b = 0$, $\left(1 - \frac{2e^{ai}-2}{2e^{ai}+2}\right)^{m-1}$, and observe that $\lim_{i \to \infty} \frac{2e^{ai}-2}{2e^{ai}+2} = 1$.

Let us then focus on the first term, which is 0 just in case

$$f(i) := -2be^{2ai+bi} + 4ae^{ai+bi} + 4ae^{ai} + 2be^{bi} = 0.$$

We show that there is a unique solution to this equation.

When $i = 0$, $f(i) = 8a$, which is positive. Observe that $f(i)$ is dominated by the $2ai + bi$ exponent, and that $\lim_{i \to \infty} f(i) = -\infty$. Thus, by the intermediate value theorem, a root exists.

We may use a generalization of Descartes' rule of signs (see Laguerre 1883) to show that $f(i)$ has at most one root. Because $2ai + bi > ai + bi > ai, bi$ (for $i \neq 0$), there is single sign change in $f(i)$ (from negative to positive). So the root

21

is unique. It may be verified (by tedious calculation of the second derivative) that this root is a maximum and that $i = 0$ is a minimum.

$\square$

Even in these special cases, no obvious analytic solution exists; $f(i) = 0$ is transcendental. Now, we may obtain an analytic solution by further constraining the equation; for example, if $a = b = 1$, then we have $-2e^{3i} + 4e^{2i} + 6e^i = 0$, for which there is a unique, closed-form solution, $i = \log 3$.

But there is nothing physically meaningful about this case (except for the fact that we have an electorate which, absent intervention, is sure to choose wrongly). Finding an optimal value of i should be estimated or solved numerically.

Although I do not see any hope of expressing the maximum of equation (6) with respect to i explicitly, numerical simulations suggest the following: First, optima are unique, with the single exception of $c_{init} = 1$.

Second, when $c_{init} \in (0.5, 1)$ and $b \gg a$, the optimum is at $i = 0$–that is, U as given by classic Condorcet is the best that the electorate can do. The intuition here is that even though this opinion leader's influence contains *something* of epistemic value, it is outweighed by the harm to group aggregation which the dependence creates. In typical cases, however, in which b is not too big relative to a, exertion of a finite amount of influence is optimal (*e.g.* Figure 4).

In contrast, third, when $c_{init} \in (0, 0.5]$, the optimum is always found at a positive but finite value of i. Because group aggregation has no epistemic value (indeed, if $c_{init} < 0.5$ it has epistemic *disvalue*), dependence is not harmful and the competence enhancement that the opinion leader provides helps–but only to a point.

It follows from these facts that it is never optimal for an opinion leader to exert influence without limit.

It might be thought that there is something idiosyncratic about the situation under consideration in light of the assumption of homogeneous competence. After all, some opinion leaders (though certainly not all) really do *know better*. When a voter is bound to an opinion leader *via* that opinion leader's influence, his vote is now more likely to be correct than it would be in the absence of dependence.

An example will show that the homogeneous competence constraint is of little relevance. Suppose that $v_1, ..., v_{2m}$ have competence c_{init}, as before, but now let v_n, the opinion leader, have competence $c^* > c_{init}$. The probability that this electorate, which includes an opinion leader who is in fact expert, will choose correctly is:[14]

$$U(i) = c^* \beta_{m|2m}(c + \rho(1 - c)) + (1 - c^*)\beta_{m+1|2m}(c - \rho c). \tag{9}$$

It is still possible, and is generally true, that the opinion leader can exert a harmful amount of influence. After all, in this case $\lim_{i \to \infty} U(i) = c^*$, which may be beaten, at least, when $i = 0$. Or, for example:

---

[14]Note that we continue to assume aggregation *via* simple majority rule here. When an expert exists, this may no longer be optimal.
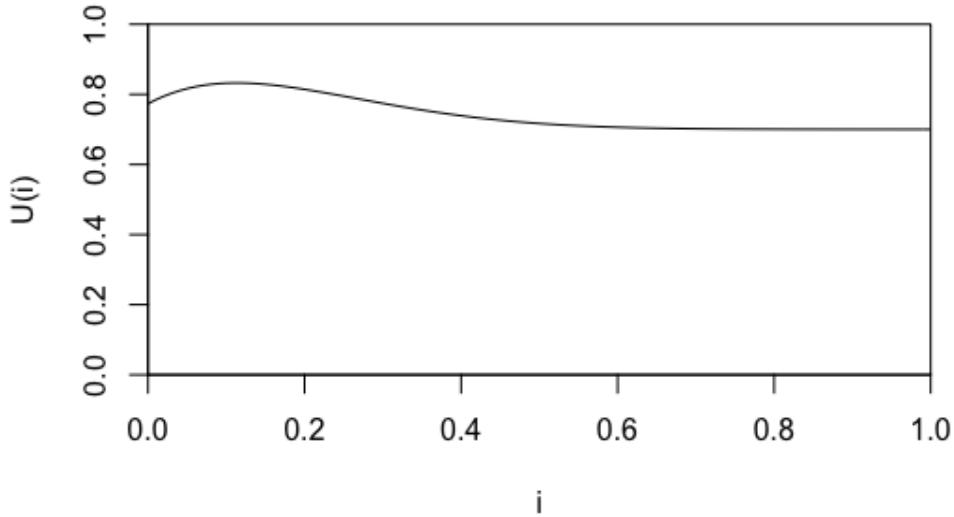
Figure 5: Effect of the expert opinion leader on the quality of the collective decision ($c^* = 0.7, c_{init} = 0.6, m = 5, a = 3, b = 3$). $U(\infty) = 0.7$; $U(0) = 0.77$; $U(0.11) = 0.83$.

4   Discussion

The model demonstrates a critical ambiguity in the exercise of political influence. It is not enough that a voter be smart, well-informed, and committed to justice. That voter must seek to maximize the efficiency of the political process, *even if that comes at the cost of his preferred political outcome*. This distinction is often elided, even as political philosophers stress the importance of political participation toward the end of justice.

Rawls, for example, says that "we are concerned with the goodness of the settled desire to take up the standpoint of justice. I assume that the members of a well-ordered society already have this desire. . . . Being the sorts of persons

they are, the members of a well-ordered society desire more than anything to act justly and fulfilling this desire is part of their good." (1999: 498). (Indeed, Rawls makes the stronger claim–that there is a responsibility to contribute to the establishment of justice "when this can be done with little cost to ourselves" (pp. 293-94).)

Similarly, John Simmons argues that "we have a duty to support and further just government, at least when this involves no great cost to ourselves" (1979: 193).

But what, exactly, does this mean? Is our duty to (1) do the best that we can to figure out what a just government looks like, and then "support and further" that understanding? Or is it to (2) "support and further" *the system* that is maximally likely to establish a just government, regardless of our judgments about the justness of the government that system creates? As we have seen, this is a critical distinction–but it is one which has gone unaddressed in the literature.

The lesson of this paper is that (2) is correct. My sense is, however, that nearly all real-world political influence comports with (1). That is, after all, how we naturally approach these matters: We develop a view of justice; we subject it to criticism; and then, once convinced of its correctness, we seek to influence the political process in accordance with it. That gets things wrong. We should set aside our own views about what is just and unjust and exert influence only insofar as it improves the accuracy of the collective judgment. In particular–at least as far as this model goes–we should avoid influence that persuades and we should pursue influence that enriches the competence of the electorate.

Conservative casino magnate Sheldon Adelson spent $83 million during the 2016 United States election cycle–a massive sum, even in our age of multi-billion-dollar elections.[15] If these considerations are correct, then Adelson may

---

[15]See https://www.opensecrets.org/overview/topindivs.php?cycle=2016&view=fc, re-

have done something wrong–immoral and perhaps irrational (if Adelson's goal was *the establishment of justice*, rather than the establishment of such-and-such because he judged it to be just). In skewing the election so strongly in favor of his own opinions, Adelson reduced the quality of the collective judgment. His influence made it less likely that justice would result, a fact that he could have and should have been aware of when contemplating how much to influence the election. (Again: The effect of his influence on the quality of the collective judgment, U, is transparent to the opinion leader.)

Of course there are real-world complexities. One could imagine Adelson reasoning as follows: "Look, George Soros pushed the political process to the left to the tune of \$22 million. Therefore, my participation, indefensible on its face, merely nullifies the deviation from optimal collective-decision making which Soros introduced. In this way, it serves the cause of justice." Maybe that's true and maybe it isn't, but if so, it suggests a prisoner's dilemma-type situation:[16] Both Adelson and Soros would prefer spending \$0 and having an efficient political process to spending millions and having an efficient political process. But since neither side trusts the other to abstain from exerting influence, the latter obtains rather than the former. Perhaps the solution to this form of political failure is for the government to proscribe campaign spending of this size.

I concede that I have not given a detailed argument for the superiority of (2) (supporting and furthering the system that is maximally likely to establish justice) over (1) (supporting and furthering X, where X is our personal under-

---

trieved 6 August 2020.

[16]Technically, the game only has a prisoner's dilemma structure if Adelson and Soros are interpreted as opinion leaders who seek justice *as they see it* rather than as the superior, process-optimizing type. The reason is that, so interpreted, each would prefer most of all to exert influence and have his opponent exert no influence, followed by refraining from influence in the face of his opponent's restraint, followed by exerting influence in the face of his opponent's exertion of influence. In any case, the dynamic described for the two relevant cases–each exerts influence and each refrains from exerting influence–holds under both interpretations.

standing of what is just). In part this is due to limitations of space; in part I think the argument is persuasive on its face. Nevertheless, there are philosophical complexities here. To mention one, note that if voting is valuable for non-epistemic reasons (*i.e.* reasons unrelated to the collective's ability to get to the truth about justice), that may mitigate things. For an opinion leader of type (1) instantiates expressive value and type (2) does not. (On expressive voting, see, *e.g.*, Brennan & Lomasky 1993.) Nevertheless, many proponents of democratic government (and its opponents, for that matter) make their arguments on epistemic grounds.[17] So the considerations raised here should be of broad interest.

Some might object that the philosophical point of this paper is paradoxical. After all, suppose that a voter believes that (1) X is the correct alternative in some political dispute; (2) if he refrains from influencing the political process, Y will be chosen rather than X; and (3) if he exerts influence X will be chosen. How can the right thing to do be that which you believe will lead to injustice?

In fact there is no paradox. There is only the appearance of paradox, a result of false views about the connection between belief and action. In fact, to act rightly, it is neither necessary nor sufficient to believe that what you are doing will produce the result you believe to be right.

As some philosophers have noted,[18] it is not sufficient because, it if were, then history's moral monsters would be exonerated. Some Nazis were simply insane, but others honestly believed, using their best judgment, that they were making the world a better place. Of course, that these Nazis acted "in good faith" in this way changes little if anything about the immorality of what they did.

The argument of this essay is that necessity does not hold either. A voter

---

[17] *E.g.* Brennan 2016, Estlund 2008, and Landemore 2013.
[18] Including Foot (2002), Schueler (2007), and Wilcox (1968).

may, perfectly coherently, find himself in the position of believing that X is just and that he should act such that Y obtains. Action gets severed from belief in this way. Because neither sufficiency nor necessity hold, the folk notions that you should "follow your conscience" or "do what you think is right" are flawed. Obeying such maxims can and often does lead people into moral error.

Finally, as mentioned in the introduction, this model suggests a solution to Wollheim's paradox (which is similar to, but different from, the case just described).[19] The paradox is as follows: Suppose one believes both that (1) we ought to implement policy X and (2) we ought to implement the policy that is chosen through a suitable democratic process. This is a common enough sentiment. Suppose, further, that such a process chooses Y rather than X (and the two are incompatible). Doesn't this put a voter in the untenable position of believing we should implement both X and *not X*?

Note, first, that the paradox only arises if (2) is interpreted *de re* rather than *de dicto*. For if it is interpreted *de dicto*, then there is no explicit contradiction, and the situation may be addressed in familiar ways (see, *e.g.*, Nelson 2019).

But if it is interpreted *de re* (or *de dicto*, for that matter), then the paradox does not arise when we think about political participation in the way suggested in this paper. Our opinion leader believes that we should implement X; he does not believe that we should implement Y. He thinks that Y is the wrong policy. So there is no inconsistency in beliefs about what we ought to do, politically.

The relevant belief, on the other side of things, is that he should *act* in a certain way (*i.e.* exert a certain amount of influence), perhaps with the understanding that this will lead to X not being realized. But, I have argued, that is perfectly consistent with having "contrary" first-order beliefs about justice.

---

[19]For critical commentary on Wollheim's paradox, see Honderich 1974, Paris & Reynolds 1978, and Weiss 1973.

5 Conclusion

A salient and troubling feature of contemporary democratic deliberation is its low quality. As there is no demand for sober and sophisticated policy debate, none is provided. Politicians, rationally enough, focus on what will get them elected: fundraising, inveighing against the idea of compromise, crafting effective propaganda (the nastier, the better), and avoiding any appearance of being vulnerable and well-rounded human beings.

But as the quality of public discourse has reached a low, interest and participation in politics remains high. While most academic work has concentrated on citizens' poor political knowledge,[20] and the effects of this ignorance on political outcomes, I have endeavored to show here that even good-willed, intelligent, and thoughtful voters may detract from important goals–goals like justice–by virtue of their political influence.

The model may be extended in various ways to accommodate contemporary political dynamics. For example: the cost to an opinion leader of exerting influence may be quantified and incorporated; we may account for heterogeneous competence throughout the electorate; and multiple sources of dependence, from multiple voters, may be analyzed. It is clear that any such extensions will have to be tackled numerically.

I conjecture that one reason for the poor political outcomes we have seen in the United States and elsewhere in recent years is an increase in dependence-enhancing political influence ("Fox News") relative to competence-enhancing influence ("disinterested educator"). If it is the case that voters, media outlets, and academics are more focused on persuading others that their own views are correct rather than enriching the electorate with their ideas, stories, and arguments, that will reduce the quality of political aggregation. Worse results

---

[20]See, *e.g.*, Althaus 2003, Brennan 2016, Converse 1964, Delli Carpini & Keeter 1996, Neuman 1986, Shenkman 2008, and Somin 2013.

are the predictable consequence.

Our collective decision-making processes have epistemic power, and we would be foolish to abandon them despite their nasty timbre. Instead, the solution to our political ills is for political actors, including voters, to think less about how they can advance justice as they see it and more about optimizing the political process. I am not sure what an optimal justice-seeking politics looks like, but I am quite sure it is not the U.S. status quo, in which politics is at its best a battle of wits and at its worst a blood sport. Our political culture should be a collegial investigation into difficult questions of import to us all. In that culture–the most effective culture for reaching moral and political truth–we are all a little more thoughtful and quiet.

6   References

Althaus, S. L. 2003. *Collective Preferences in Democratic Politics: Opinion Surveys and the Will of the People*. Cambridge: Cambridge University Press.

Austin-Smith, D. 1990. Information transmission in debate. *American Journal of Political Science* 34: 124-52.

Bahadur, R. R. 1961. A representation of the joint distribution of responses to n dichotomous items. In *Studies in Item Analysis and Prediction*, ed. H. Solomon, 158-68. Stanford: Stanford University Press.

Berend, D. & Sapir, L. 2005. Monotonicity in Condorcet Jury Theorem. *Social Choice and Welfare* 24: 83-92.

––. 2007. Monotonicity in Condorcet's Jury Theorem with dependent voters. *Social Choice and Welfare* 28: 507-28.

Berg, S. 1993. Condorcet's Jury Theorem, dependency among voters. *Social Choice and Welfare* 10: 87-95.

––. 1996. Condorcet's Jury Theorem and the reliability of majority voting.

Group Decision and Negotiation 5: 229-38.

Bergmann, M. 2009. Rational disagreement after full disclosure. *Episteme* 6: 336-53.

Boland, P. J. 1989. Majority systems and the Condorcet Jury Theorem. *The Statistician* 38: 181-89.

Boland, P. J., Proschan, F., & Tong, Y. L. 1989. Modelling dependence in simple and indirect majority systems. *Journal of Applied Probability* 26: 81-89.

Brennan, G. & Lomasky, L. 1993. *Democracy & Decision: The Pure Theory of Electoral Preference*. New York: Cambridge University Press.

Brennan, J. 2016. *Against Democracy*. Princeton: Princeton University Press.

Christensen, D. 2007. Epistemology of disagreement: the good news. *Philosophical Review* 119: 187-217.

——. 2009. Disagreement as evidence: the epistemology of controversy. *Philosophy Compass* 4: 756-67.

Condorcet. 1785. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Paris: Imprimerie Royale.

Converse, P. E. 1964. The nature of belief systems in mass publics. In *Ideology and Discontent*, ed. D. E. Apter, 164-93. New York: Free Press.

Delli Carpini, M. X. and Keeter, S. 1996. *What Americans Know about Politics and Why It Matters*. New Haven: Yale University Press.

Dietrich, F. 2008. The premises of Condorcet's Jury Theorem are not simultaneously justified. *Episteme* 5: 56-73.

Dietrich, F. & List, C. 2004. A model of jury decision where all jurors have the same evidence. *Synthese* 142: 175-202.

Dietrich, F. & Spiekermann, K. 2013a. Epistemic democracy with defensible premises. *Economics and Philosophy* 89: 87-120.

——. 2013b. Independent opinions? On the causal foundations of belief forma-

tion and jury theorems. *Mind* 122: 655-85.

——. 2020. Jury theorems. In *The Routledge Handbook of Social Epistemology*, eds. M. Fricker, P. J. Graham, D. Henderson, & N. J. L. L. Pedersen, 386-96. New York: Routledge.

Elga, A. 2007. Reflection and disagreement. *Noûs* 41: 478-502.

Estlund, D. 2008. *Democratic Authority: A Philosophical Framework*. Princeton: Princeton University Press.

Ewin, R. E. 1967. Wollheim's paradox of democracy. *Australasian Journal of Philosophy* 45: 356-57.

Feddersen, T. & Pesendorfer, W. 1996. The swing voter's curse. *American Economic Review* 86: 408-24.

——. 1997. Voting behavior and information aggregation in elections with private information. *Econometrica* 65: 1029-58.

——. 1999. Abstention in elections with asymmetric information and diverse preferences. *American Political Science Review* 93: 381-98.

Feldman, R. 2007. Reasonable religious disagreements. In *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, ed. L. M. Antony, 194-214. New York: Oxford University Press.

Foot, P. 2002. Moral relativism. In *Moral Dilemmas: And Other Topics in Moral Philosophy*, ed. P. Foot, 20-36. New York: Oxford University Press.

Goldman, A. 1999. *Knowledge in a Social World*. Oxford: Oxford University Press.

Grofman, B. & Feld, S. L. 1983. Determining optimal weights for expert judgment. In *Information Pooling and Group Decision Making: Proceedings of the Second University of California, Irvine, Conference on Political Economy*, eds. B. Grofman & G. Owen, 167-72. Greenwich, CT: JAI Press.

Grofman, B., Owen, G., & Feld, S. L. 1983. Thirteen theorems in search of the

truth. *Theory and Decision* 15: 261-78.

Honderich, T. 1974. A difficulty with democracy. *Philosophy & Public Affairs* 3: 221-26.

Hummel, P. 2010. Jury theorems with multiple alternatives. *Social Choice and Welfare* 34: 65-103.

Kaniovski, S. 2008a. The exact bias of the Banzhaf measure of power when votes are neither equiprobable nor independent. *Social Choice and Welfare* 31: 281-300.

——. 2008b. Straffin meets Condorcet: What can a voting power theorist learn from a jury theorist? *Homo Oeconomicus* 25: 181-202.

——. 2010. Aggregation of correlated votes and Condorcet's Jury Theorem. *Theory and Decision* 69: 43-68.

Kaniovski, S. & Zaigraev, A. 2011. Optimal jury design for homogenous juries with correlated votes. *Theory and Decision* 71: 439-59.

Kelly, T. 2005. The epistemic significance of disagreement. In *Oxford Studies in Epistemology, Volume 1*, eds. T. S. Gendler & J. Hawthorne, 167-96. Oxford: Oxford University Press.

Ladha, K. K. 1992. The Condorcet Jury Theorem, free speech, and correlated votes. *American Journal of Political Science* 36: 617-34.

——. 1993. Condorcet's Jury Theorem in light of de Finetti's theorem. *Social Choice and Welfare* 10: 69-85

——. 1995. Information pooling through majority-rule voting: Condorcet's Jury Theorem with correlated votes. *Journal of Economic Behavior and Organization* 26: 353-72.

Laguerre, M. 1883. Mémoire sur la théorie des équations numériques. *Journal de Mathématiques Pures et Appliquées* 9: 99-146.

Lam, L. & Suen, C. Y. 1996. Majority vote of even and odd experts in a poly-

chotomous choice situation. *Theory and Decision* 41: 13-36.

Landemore, H. 2013. *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many.* Princeton: Princeton University Press.

Lazarsfeld, P. F. 1956. *Some observations on dichotomous systems.* New York: Columbia University Sociology Department.

List, C. & Goodin, R. E. 2001. Epistemic democracy: generalizing the Condorcet Jury Theorem. *Journal of Political Philosophy* 9: 277-306.

List, C. & Pettit, P. 2004. An epistemic free-riding problem? In *Karl Popper: Critical Appraisals*, eds. P. Catton & G. Macdonald, 128-58. Abingdon: Routledge.

List, C. & Spiekermann, K. 2016. The Condorcet Jury Theorem and voter-specific truth. In *Goldman and His Critics*, eds. B. P. McLaughlin & H. Kornblith, 219-31. Malden, MA: John Wiley & Sons.

Miller, N. R. 1996. Information, individual errors, and collective performance: empirical evidence on the Condorcet Jury Theorem. *Group Decision and Negotiation* 5: 211-28.

Mulligan, T. 2018a. *Justice and the Meritocratic State.* New York: Routledge.

——. 2018b. Plural voting for the twenty-first century. *The Philosophical Quarterly* 68: 286-306.

——. 2021. The epistemology of disagreement: Why not Bayesianism? *Episteme* 18: 587-602.

Nelson, M. 2019. Propositional attitude reports. In the *Stanford Encyclopedia of Philosophy (Spring 2019 Edition)*, ed. E. N. Zalta, URL = `https://plato.stanford.edu/archives/spr2019/entries/prop-attitude-reports/`.

Neuman, W. R. 1986. *The Paradox of Mass Politics: Knowledge and Opinion in the American Electorate.* Cambridge: Harvard University Press.

Nitzan, S. & Paroush, J. 1982. Optimal decision rules in uncertain dichotomous

choice situations. *International Economic Review* 23: 289-97.

——. 1984. The significance of independent decisions under uncertain dichotomous choice situations. *Theory and Decision* 17: 47-60.

——. 1985. *Collective Decision Making: An Economic Outlook*. Cambridge: Cambridge University Press.

Paris, D. C. & Reynolds, J. F. 1978. Paradox, rationality, and politics: Wollheim's democracy. *Journal of Politics* 40: 956-83.

Paroush, J. 1990. Multi-choice problems and the essential order among decision rules. *Economics Letters* 32: 121-25.

——. 1998. Stay away from fair coins: a Condorcet Jury Theorem. *Social Choice and Welfare* 15: 15-20.

Peleg, B. & Zamir, S. 2012. Extending the Condorcet Jury Theorem to a general dependent jury. *Social Choice and Welfare* 39: 91-125.

Rawls, J. 1999. *A Theory of Justice: Revised Edition*. Cambridge: Harvard University Press.

Reichenbach, H. 1956. *The Direction of Time*. Berkeley: University of California Press.

Schueler, G. F. 2007. Is it possible to follow one's conscience? *American Philosophical Quarterly* 44: 51-60.

Shapley, L. & Grofman, B. 1984. Optimizing group judgmental accuracy in the presence of interdependencies. *Public Choice* 43: 329-43.

Shenkman, R. 2008. *Just How Stupid Are We? Facing the Truth about the American Voter*. New York: Basic Books.

Simmons, A. J. 1979. *Moral Principles and Political Obligations*. Princeton: Princeton University Press.

Somin, I. 2013. *Democracy and Political Ignorance: Why Smaller Government Is Better*. Stanford: Stanford University Press.

Van Der Geest, P. A. G. 2005. The binomial distribution with dependent Bernoulli trials. *Journal of Statistical Computation and Simulation* 75: 141-54.

Van Inwagen, P. 2010. We're right. They're wrong. In *Disagreement*, eds. R. Feldman & T. A. Warfield, 10-28. New York: Oxford University Press.

Weiss, D. D. 1973. Wollheim's paradox: survey and solution. *Political Theory* 1: 154-70.

Wilcox, J. T. 1968. Is it always right to do what you think is right? *Journal of Value Inquiry* 2: 95-107.

Wollheim, R. 1962. A paradox in the theory of democracy. In *Philosophy, Politics and Society*, eds. P. Laslett & W. G. Runciman, 71-87. New York: Barnes and Noble.

Zaigraev, A. & Kaniovski, S. 2013. A note on the probability of at least k successes in n correlated binary trials. *Operations Research Letters* 41: 116-20.