

Jury Theorems

Franz Dietrich
Paris School of Economics & CNRS

Kai Spiekermann
London School of Economics

Working Paper
1 August 2016

Abstract

We give a review and critique of jury theorems from a social-epistemology perspective, covering Condorcet's (1785) classic theorem and several later refinements and departures. We assess the plausibility of the conclusions and premises featuring in jury theorems and evaluate the potential of such theorems to serve as formal arguments for the 'wisdom of crowds'. In particular, we argue (i) that there is a fundamental tension between voters' independence and voters' competence, hence between the two premises of most jury theorems; (ii) that the (asymptotic) conclusion that 'huge groups are infallible', reached by many jury theorems, is an artifact of unjustified premises; and (iii) that the (non-asymptotic) conclusion that 'larger groups are more reliable', also reached by many jury theorems, is not an artifact and should be regarded as the more adequate formal rendition of the 'wisdom of crowds'.

1 Introduction

Jury theorems form the technical core of arguments for the 'wisdom of crowds', the idea that large democratic decision-making bodies outperform small undemocratic ones when it comes to identifying factually correct alternatives. The popularity of jury theorems has spread across various disciplines such as economics, political science, philosophy, and computer science. A 'jury theorem' is a mathematical theorem about the probability of correctness of majority decisions between two alternatives. The existence of an objectively correct (right, better) alternative is the main metaphysical assumption underlying jury theorems. This involves an epistemic, outcome-based rather than purely procedural conception of democracy: the goal of democratic decision-making is to 'track the truth', not to fairly represent

people's views or preferences (Cohen 1986). Typical jury theorems conclude that 'crowds are wise' in one or both of two senses:

The growing-reliability thesis: Larger groups are better truth-trackers. That is, they are more likely to select the correct alternative (by majority) than smaller groups or single individuals.

The infallibility thesis: Huge groups are infallible truth-trackers. That is, the likelihood of a correct (majority) decision tends to full certainty as the group becomes larger and larger.

Jury theorems differ considerably in their premises (axioms) about voters. They often rest on two premises, an 'independence' axiom and a 'competence' axiom, each of which may take various forms. For instance, the first of all jury theorems, attributable to the French enlightenment philosopher and mathematician Nicolas Marquis de Condorcet (1785), concludes that both of the aforementioned theses hold, based on particularly simple premises:

Condorcet's independence premise: The voters have independent probabilities of voting for the correct alternative.

Condorcet's competence premise: These probabilities exceed $1/2$, and are the same for all voters.

Following our analysis, the infallibility thesis emerges as incorrect in almost all real applications. Worse, this thesis does not even seem helpful as an approximation, idealization or paradigm of how large-scale democracy performs. It is fair to say that those classical jury theorems which conclude that huge groups are infallible – however beautiful they might be – have played a misleading role as a model of democratic decision making. Their overly optimistic conclusion has led the debate astray, suggesting to some that the infallibility thesis might be true after all, while suggesting to others that 'something' must be wrong with jury-theorem-based arguments in general. Neither reaction is justified. We shall (i) pinpoint what goes wrong in (the premises of) some naive jury theorems, and (ii) show how other jury theorems avoid flawed premises. Non-naive jury theorems reach the growing-reliability conclusion, but not the infallibility conclusion. This suggests that the growing-reliability thesis is the more appropriate formal rendition of the wisdom of crowds. That thesis, by itself, gives strong epistemic support for (majoritarian) democracy. The infallibility thesis would have given additional support – but it is not tenable, and should be taken off the agenda after having haunted the literature for decades.

We shall give a selective review of jury theorems and our own critical assessment of their suitability for formal arguments for the ‘wisdom of crowds’. We begin with a naive Condorcetian jury theorem (Section 2), which we then gradually refine into jury theorems with more plausible premises (Sections 3–5). We then discuss further jury theorems (Section 6), key objections (Section 7), and strategic voting (Section 8), before offering a concluding assessment (Section 9).

2 A naive jury theorem

We consider a group of individuals deciding by majority vote between two alternatives, such as to convict or acquit a defendant, or to keep or abolish a law. To be able to vary the group size, we consider an infinite reservoir of individuals labelled $i = 1, 2, \dots$ and take the group of size n to consist of the first n individuals $1, \dots, n$. Each individual votes for exactly one alternative. The alternative receiving more votes wins. To avoid ties, the group size n is throughout an *odd* number: $n \in \{1, 3, 5, \dots\}$.¹ Exactly one of the alternatives is ‘correct’ or ‘better’ in an objective, voter-independent sense; it is called the unknown *state (of the world)*.

We first state the jury theorem in a simple and common (yet as we shall see problematic) version. The only model ingredients are events R_1, R_2, \dots representing correct voting by individuals $1, 2, \dots$, respectively. Alternatively, the model ingredients could be random variables $\mathbf{v}_1, \mathbf{v}_2, \dots$ representing the votes of individuals $1, 2, \dots$ and another random variable representing the true state, all ranging over the same binary set of alternatives, e.g., the set {‘convict’, ‘acquit’}, or {‘abolish’, ‘keep’}, or $\{0, 1\}$; each correctness event R_i is then defined as the event that \mathbf{v}_i coincides with the state.²

We are ready to state the simple Condorcetian jury theorem, beginning with its two axioms (e.g., Grofman et al. 1983).

Unconditional independence (UI): The correctness events R_1, R_2, \dots are (unconditionally) independent.

Unconditional competence (UC): The (unconditional) correctness probability $p = P(R_i)$, the (*unconditional*) *competence*, (i) exceeds $1/2$ and (ii) is the same for each voter i .

¹One could work without excluding even n , by assuming that ties are broken using a fair coin.

²Events and random variables are of course defined relative to a background probability space, i.e., a set of possible worlds Ω , a notion of ‘events’ (one could count all subsets of Ω as events, or, more generally, all subsets in a given σ -algebra of subsets), and a probability function P defined on the set of events. We shall never mention the probability space explicitly as a model ingredient.

Theorem 1. *Assume UI and UC. As the group size increases, the probability of a correct majority³ (i) increases⁴ (growing reliability), and (ii) tends to one (infallibility).*

Mathematically, the infallibility conclusion is an easy consequence of the law of large numbers, which implies that, under UI and UC, as the group size tends to infinity, the correctness proportion converges to the correctness probability $p = P(R_i)$ (with probability one), so that the probability of a correct majority tends to one. The growing-reliability conclusion is harder to prove.⁵

3 A state-sensitive jury theorem

The independence assumption UI is highly problematic, even if voters do not communicate with each other. Why? This section explains *one* of the problems, and presents a jury theorem that fixes it (other problems are addressed in Sections 4 and 5). Binary decision problems often display an asymmetry between the alternatives: one alternative is simpler to identify as correct than the other one. Guilt of a defendant might be easier to detect than innocence, or vice versa; global warming might be easier to detect than its absence, or vice versa; and so on. This sort of truth-tracking asymmetry is the rule, not the exception. It renders the event R_i that individual i identifies the state positively correlated with the event of the simpler-to-identify state, because correct voting is more likely given the simpler-to-identify state than given the harder-to-identify state. For example, if guilt is simpler to identify than innocence, a juror is more likely to get it right given guilt than given innocence. As all correctness events R_1, R_2, \dots correlate positively with the same event (of the simpler-to-identify state), they normally correlate positively *with one another*.⁶ So UI is violated.

This problem of correlation ‘via’ the state can be avoided by holding the state fixed, that means, conditionalizing on the state. We first enrich Section 2’s model by another ingredient: a random variable \mathbf{x} , interpreted in this section as the state, i.e.,

³This is the probability that the number of individuals $i \in \{1, \dots, n\}$ such that R_i obtains exceeds $n/2$.

⁴Throughout, ‘increases’ is used in its weak sense. In fact, the increase in Theorem 1 is strict, provided the correctness events R_i do not have probability one. In the other theorems stated the increase is also strict, except for extreme cases.

⁵The proof relies on a recursive formula for the probability of a correct majority as a function of n . This formula is stated in Grofman et al. (1983), but its proof is hard to find (see however Dietrich and Spiekermann 2013a, Step 2 in Appendix C).

⁶We say ‘normally’ rather than ‘necessarily’ since there are some far-fetched mathematical counterexamples.

the correct alternative.⁷ Formally speaking, nothing hinges on this interpretation of \mathbf{x} which could be almost any random variable.⁸

We can now revise the axioms and theorem as follows.

Conditional Independence (CI): The correctness events R_1, R_2, \dots (or equivalently the votes $\mathbf{v}_1, \mathbf{v}_2, \dots$)⁹ are independent conditional on any value of \mathbf{x} .

Conditional Competence (CC): For any value x of \mathbf{x} , the conditional correctness probability $p^x = P(R_i|x)$, the *competence (conditional on x)*, (i) exceeds 1/2 and (ii) is the same for all voters i (but may vary with x).

Theorem 2. *Assume CI and CC. As the group size increases, the probability of a correct majority (i) increases (growing reliability), and (ii) tends to one (infallibility).*

This jury theorem reaches the same conclusions as Theorem 1, but on the basis of ‘state-conditional’ axioms. The state x with higher competence p^x (if it exists) is the easier-to-identify state discussed at the start of the section.

4 The fundamental tension between independence and competence

Different votes can be correlated via the objective state they track – a problem solved above by working with the independence axiom CI rather than UI, thereby fixing the state and blocking the correlation via the state. Unfortunately, votes can also be correlated via several other circumstances. So fixing just the state – the most common form of conditionalization in the literature – does not yet secure independence (Dietrich and List 2004, Dietrich 2008, Dietrich and Spiekermann 2013a; see also Ladha 1992, 1995). *Any common cause* of votes is a potential

⁷Recall that Section 2’s initial model has as ingredients the correctness events R_1, R_2, \dots , or alternatively the votes $\mathbf{v}_1, \mathbf{v}_2, \dots$ and a state variable (with R_i then defined as the event that \mathbf{v}_i matches the state). If we adopt the alternative ingredients and if the new ingredient \mathbf{x} is itself the state (following our current interpretation), then the initial model already contains \mathbf{x} and need not be extended.

⁸ \mathbf{x} could take any sort of values, with the only constraining assumption that each value has positive probability in order to render conditionalization meaningful. This implies that \mathbf{x} is a discrete random variable, i.e., takes only finitely or countably infinitely many values. (Everything could be generalized to possibly continuous \mathbf{x} .)

⁹The equivalent formulation in terms of votes assumes that \mathbf{x} is or more generally subsumes (determines) the state, an assumption in line with our current and later interpretations of \mathbf{x} . Under this assumption, the equivalence holds since R_i and \mathbf{v}_i are then interdefinable given \mathbf{x} : R_i holds if and only if \mathbf{v}_i matches the state determined by \mathbf{x} .

source of dependence. Consider common evidence, such as, in a court case, witness reports and the defendant’s facial expression, or, among scientists, experimental data. Evidence may or may not support the truth: it may be truth-conducive or misleading. For instance, the defendant’s friendly facial expression presumably supports innocence and is misleading in case of guilt. Plausibly, the correctness events R_1, R_2, \dots correlate positively with the event of truth-conducive (i.e., non-misleading) common evidence, and thereby correlate positively *with one another*. Voters can also be influenced by common causes that are *non-evidential* such as distracting heat: such causes lack an objective bearing on the true state, and yet they influence people’s epistemic performance and thereby threaten independence. Jurors are more likely to vote well in agreeable room temperature; votes are thus correlated via room temperature.¹⁰

The strategy to restore independence should by now be familiar: one should conditionalize on the common causes of votes. So we now reinterpret the variable \mathbf{x} on which we conditionalize in axioms as representing not just the state, but in addition all common causes of votes, the ‘circumstances’. In the terminology of Dietrich (2008) and Dietrich and Spiekermann (2013a), \mathbf{x} represents the specific *decision problem* faced by the group. For such \mathbf{x} , axiom CI becomes plausible. Have we thus rehabilitated Theorem 2 as a formal argument for the ‘wisdom of crowds’? Unfortunately not, because our rich interpretation of \mathbf{x} renders the competence axiom CC implausible. Why?

Generally, whether a voter is competent – i.e., more often right than wrong – depends on the reference class considered. Plausibly, a voter is more often right than wrong among all conceivable yes/no questions, or all guilty/innocent questions. While competent within such an all-encompassing reference class, a voter is presumably not competent within a reference class in which certain misleading evidence is always present, such as all guilty/innocent questions where the defendant is guilty even though he has an alibi. Once we conditionalize on the full decision problem, we fine-grain the reference classes and effectively randomize only over parameters other than the state, common evidence, and other common causes.¹¹ Just imagine a decision problem (a value of \mathbf{x}) characterized by severely misleading evidence, say a decision problem in which an innocent defendant unluckily looks exactly like the true murderer captured on CCTV. In this reference class a juror will be incompetent: within it he will get it right only rarely, e.g., when inattentive to (misleading)

¹⁰Common causes, whether evidential or non-evidential, threaten not just unconditional independence UI but also state-conditional independence since the correlations do not disappear by conditionalizing on the state.

¹¹What can still vary is, say, the voter’s level of awakesness, attention, hunger, or back pain, all of which may indeed influence the voter’s truth-tracking ability and hence his correctness probability.

evidence. Most court cases (or decision problems) are not of this unfortunate kind: most have mainly truth-conducive evidence. So the voter is more often right than wrong across a wider cross-problem reference class. This observation is, however, irrelevant for the problem-conditional competence axiom CC, which conditionalizes on a specific decision problem rather than ‘averaging out’ the unlucky cases of misleading evidence.

Ironically, the *problem-conditional* notion of probability renders independence (CI) defensible but competence (CC) unjustified, whereas a cross-problem-randomizing notion of probability – whether Section 2’s *unconditional* or Section 3’s *state-conditional* notion – renders competence (UC or CC) more justified but independence (UI or CI) implausible. So, Theorems 1’s and 2’s premises are not jointly justifiable, regardless of how much we conditionalize on or ‘pack’ into \mathbf{x} . The following table summarizes this dilemma.¹²

	unconditional	state-conditional	problem-conditional
independence	UI	CI for $\mathbf{x} = \textit{state}$	CI for $\mathbf{x} = \textit{problem}$
axiom	implausible	implausible	plausible
competence	UC	CC for $\mathbf{x} = \textit{state}$	CC for $\mathbf{x} = \textit{problem}$
axiom	plausible in homog. groups	plausible in homog. groups	implausible

5 A problem-sensitive jury theorem

Following the previous section, we interpret the variable \mathbf{x} as capturing the group’s specific decision problem, including the common causes of voters. So the independence axiom CI is plausible, but the competence axiom CC is untenable. It is tempting to replace CC by the unconditional competence axiom UC, but unfortunately the combination of CI and UC – two potentially justified premises – does not lend itself to a jury theorem: it does not imply the growing-reliability thesis. We therefore weaken CC to a more plausible axiom: *tendency to competence*. Stating this axiom

¹²An altogether different approach avoids the need of conditionalizing on the decision problem and defends the independence axiom in its unconditional form UI by interpreting probabilities differently in the first place: the probabilities captured by the function P are now taken to already incorporate (be ‘posterior on’) the decision problem. The state and the common causes are thus fixed rather than drawn randomly. This interpretation of probability turns Theorem 1 into Dietrich’s (2008) fixed-problem jury theorem. Theorem 1’s independence axiom UI is now justified, but its competence axiom UC is no longer justified. The problem is not that UC *must* fail, but that it is *unknown* whether it holds: since the fixed state is unknown, so is the question of whether the fixed problem has truth-conducive or misleading circumstances. So the alternative interpretation of probability also fails to render Theorem 1’s premises jointly justified. We shall not adopt this interpretation.




Figure 1: Example of Tendency to Competence.

requires a short preparation. A voter i 's problem-specific competence $p^{\mathbf{x}} = P(R_i|\mathbf{x})$ depends on the (randomly drawn) problem \mathbf{x} , and is thus itself a random variable. Its value is above $1/2$ for 'easy' problems (values of \mathbf{x}), and below $1/2$ for 'difficult' problems with misleading evidence or other epistemically harmful circumstances. A discrete real-valued random variable, in our case the competence variable, *tends to exceed* $1/2$ if the value $1/2 + \epsilon$ is at least as probable as the symmetrically opposed value $1/2 - \epsilon$, and this for all $\epsilon > 0$. An illustration is given in Figure 1.¹³ Here the competence level is for, instance, more likely to be $0.7 = 0.5 + 0.2$ than to be $0.3 = 0.5 - 0.2$. This and all other symmetrical comparisons are indicated by dashed lines.

Tendency to Competence (TC): A voter i 's competence $p^{\mathbf{x}} = P(R_i|\mathbf{x})$ (as a function of \mathbf{x}) (i) tends to exceed $1/2$, and (ii) is the same for all voters i .

Axiom TC weakens CC: it retains CC's homogeneity part, but weakens CC's first part by allowing voters to sometimes be incompetent. By using TC rather than the implausible axiom of CC, we no longer reach the implausible infallibility conclusion, while retaining the growing-reliability conclusion (Dietrich and Spiekermann 2013a):

Theorem 3. *Assume CI and TC. As the group size increases, the probability of a correct majority (i) increases (growing reliability), and (ii) tends to a value which is below 1 (no infallibility) unless CC holds.*¹⁴

¹³We thank the developers of the open source matplotlib library.

¹⁴To disambiguate, 'unless' means 'if and only if it is *not* the case that'. See Dietrich and Spieker-

This theorem gives group deliberation and communication a new role. Classical jury theorems suggest that deliberation might be harmful by threatening independence (e.g. Anderson 2006, Rawls 1971, pp. 314-5). But deliberation does not undermine the new problem-conditional independence axiom CI: insofar as deliberation leads to information exchange and hence to additional common evidence, this common evidence is incorporated into the decision problem, so that common-evidence-caused correlations are automatically ‘conditionalized away’. Instead, deliberation is beneficial: it ideally renders voters more competent¹⁵ and thereby the group’s majority judgment more reliable.

6 Further jury theorems

In this section we provide a short, non-exhaustive overview of other jury theorems (setting aside the sort of concerns raised above, although they still apply). The most frequent starting point is the classic Theorem 1, in which either axiom could be weakened.

Weakening axiom UC by simply dropping its homogeneity condition, hence allowing voter-specific competence $p_i = P(R_i)$, has dramatic consequences: neither the growing-reliability nor the infallibility thesis still follows. Instead the probability of a correct majority can *decrease* in group size and converge to 1/2, so that huge groups are as bad as a fair coin. This happens if the sequence of competence levels p_1, p_2, \dots strictly decreases towards 1/2, so rapidly that newly added voters are much less competent than existing voters and thereby pull the majority’s reliability down (Paroush 1998). Some restriction on *how* competence varies across voters is thus needed for any ‘wisdom of crowds’ conclusion. One could weaken UC either to the condition that all p_i exceed $1/2 + \epsilon$ for a voter-independent $\epsilon > 0$ (Paroush 1998), or to the condition that average competence $\frac{1}{n} \sum_{i=1}^n p_i$ converges to a value above 1/2 (Dietrich 2008; for related or more general results, see Berend and Paroush 1998 and Owen et al. 1989). This preserves the infallibility conclusion, but not the growing-reliability conclusion – the ‘wrong’ conclusion is preserved, one might complain. Alternatively, one no longer assumes fixed identity of voters but draws the group of any given size n randomly from a given huge (finite) pool of potential voters i with competence levels $p_i > 1/2$. The probability of a correct majority vote

mann (2013a) for the generalization to possibly non-discrete \mathbf{x} , and for a version of the result which concludes that the group’s performance increases *strictly* (something achieved by strengthening TC through using a strict notion of ‘tendency to exceed 1/2’).

¹⁵A voter’s problem-specific competence $p^{\mathbf{x}}$ presumably tends to exceed 1/2 *more strongly* for a ‘post-deliberation decision problem’ \mathbf{x} enriched by new common evidence than for a ‘no-deliberation decision problem’ \mathbf{x} with less or no common evidence.

is then doubly random: the identity of the voters and their votes are random. The growing-reliability conclusion is then restored (Berend and Sapir 2005).

The other candidate for modification is the independence axiom. Boland (1989) and Boland et al. (1989), for example, discuss the influence of an opinion leader. Before the opinion leader is consulted, voters' judgments obey UI and UC. Each voter (other than the opinion leader) has the same independent probability of following the opinion leader in his vote; votes thus violate UI. The infallibility conclusion is still reached if (and only if) the probability of following is not too high, specifically below $1 - \frac{1}{2p}$ (see also Goodin and Spiekermann 2012). There are many other ways to adapt or weaken independence. For instance, some jury theorems assume votes to be interchangeable rather than independent, giving up the infallibility conclusion (Ladha 1993, Dietrich and Spiekermann 2013a); others are based on causal networks, again giving up infallibility (Dietrich and List 2004, Dietrich and Spiekermann 2013b); Kaniovski (2010) analyses the group's reliability as a function of the dependence structure among voters; and Pivato (2016) shows that the infallibility conclusion is still reachable under surprisingly strong forms of dependence, using suitable voting rules.

List and Goodin (2001) importantly generalize Theorem 2's infallibility conclusion to the case of $k \geq 2$ options, still assuming one option is correct. They generalize CC: given the correct option x , a voter is more likely to vote for x than for y , and this for every incorrect option y . Interestingly, the probability of correct voting no longer needs to exceed $1/2$ for $k > 2$ options. It can be conjectured that the growing-reliability conclusion also generalizes to $k \geq 2$ options. The upshot is that jury theorems and 'wisdom of crowds' arguments need not be restricted to binary decision problems.

Going beyond majority voting, some work asks which voting rule is epistemically optimal, where 'optimality' could be cashed out differently, e.g., by maximizing the probability of a correct outcome (e.g., Nitzan and Paroush 1982, Shapley and Grofman 1984, Ben-Yashar and Nitzan 1997, Dietrich 2006, Pivato 2013). The generic finding is that, under an independence axiom of type UI or CI, the optimal rule is a weighted super-, sub- or simple majority rule in which a voter's weight is a well-defined function of his competence (and becomes negative in case of incompetence). Simple majority rule is optimal in case of equal competence and symmetric options.

7 Objections and replies

One objection is David Estlund's 'disjunction problem' (Estlund 2008, p. 228-30). We give a new reconstruction of the problem, showing its appeal but also its limits.

In a choice, say between building a roadbridge or a footbridge, each option is a disjunction of several finer sub-options. For instance, a roadbridge might have one, two or more lanes, and might or might not have street lights and plants. Alternatively, the voters could decide between a *specific* roadbridge (say one with 2 lanes, no street lights, but plants) or any other road- or footbridge. The first option has become finer, the second coarser.¹⁶ The change in options may considerably alter a voter's pair of state-conditional competence levels (the correct-voting probabilities given correctness of the first or second option). This by itself is not a problem, but just a feature, we maintain. What *is* a problem is the following. If one option is much more specific than the other, then a voter's conditional competence may easily be below 1/2 given the specific state of the world (and close to 1 given the unspecific state). For correctness of a highly specific option is often subjectively unlikely, hence not recognized – just as it seems unlikely that you ate exactly 139 corn flakes this morning, even if you did. This argument threatens the plausibility of the competence axiom CC, but only for decision problems whose options are highly unequal in specificity. Also, the competence axioms UC and TC remain plausible.¹⁷ Since Estlund neither distinguishes between conditional and unconditional competence nor explicitly restricts the scope of the problem, his objection is overstated.

The disjunction problem is not the only objection raised against competence axioms. Some scholars question competence on empirical grounds. There is, however, one theoretical argument against systematic incompetence: systematically worse-than-random judgments are *unstable*, since once a voter becomes aware of his incompetence he can normally achieve competence by simply reversing all of his judgments (Dietrich 2008, Goodin and Spiekermann 2012).

Another frequent objection is that jury theorems invoke variables whose exact nature is unknown in practice. One usually knows neither exact competence levels nor the precise nature and distribution of the decision problem \mathbf{x} in Theorem 3, as the precise common causes are unknown. However, accepting a competence axiom (UC, CC or TC) only requires ascertaining some inequalities, without having to know the exact behaviour of the variables involved.

Finally, on the most fundamental level, some critics suggest that there often is no truth to be tracked (e.g. Miller 1992, p. 56, Black 1958, p. 163). For instance, it is sometimes claimed that there is no objective truth in political choices, which are supposedly not about facts but about values or preferences. Whether this is so

¹⁶Presumably the new first option, the specific roadbridge, can be the correct option only if the old first option, the unspecific roadbridge, was the correct option.

¹⁷Regarding UC, a voter is unconditionally likely to vote correctly since he is likely to believe the unspecific option is correct and since that option is itself likely to be correct. It is only conditional on correctness of the specific option that the voter is likely to get it wrong.

arguably depends on how the question is asked (Landemore 2013, ch. 8). Asking British voters whether they “prefer” the United Kingdom (UK) to remain in the European Union (EU) suggests a non-epistemic preference elicitation. Asking whether EU membership is “better” or “better for the UK” seems to ask two (different) epistemically framed questions. In fact, the question in the EU referendum on 23 June 2016 was whether the UK “should” remain in or leave the EU. Whether this is an epistemic issue, i.e., whether there is a fact about what the UK “should” do, depends on how one interprets “should” and on meta-ethical commitments.

An interesting hybrid view is that, while there *is* an objective fact, it is a group-specific fact about which option is *preferred* by the majority, i.e., the “will of the majority”. The question then is: does a voter vote according to his own preference, or according to his belief about the majority preference? In the latter case, the standard jury-theorem setup directly applies. In the former case, majority outcomes may seem to be correct by definition. But, on a more sophisticated version of the view, a voter can be mistaken about his own true preference, so that his vote tracks a *voter-specific truth* about his preference. Under plausible conditions, large groups are likely to track the majority preference (Miller 1986, Goldman 1999, p. 323ff, Goodin and Spiekermann 2015, List and Spiekermann forthcoming).

8 Strategic voting

It is well-known that conflicts of interests between voters may lead to strategic incentives. Surprisingly, strategic voting may also occur in purely epistemic contexts where voters share the common goal of objectively correct group decisions; this insight has sparked a large and active literature, particularly among economists (Austen-Smith and Banks 1996, Feddersen and Pesendorfer 1999, Peleg and Zamir 2012, Bozbay, Peters and Dietrich 2014, among many others). We only sketch the basic idea. Suppose you are one of 11 jurors voting on whether to convict the defendant. Your private information suggests guilt. Should you vote ‘convict’? Strategically, your vote should be chosen on the assumption that it makes a difference (is ‘pivotal’); so that 5 other jurors vote ‘convict’ and 5 vote ‘acquit’. In any other situation your vote would be irrelevant for the outcome and can therefore be ignored. Hence, after adding the hypothetical information about other voters to your own information, ‘convict’ is supported 6 times, and ‘acquit’ 5 times. This is almost an informational tie, which (let’s assume) does not justify conviction. You should vote ‘acquit’, although your private information alone suggests conviction. Paradoxically, if every voter reasons like this, all jurors vote to acquit even when they all hold private information suggesting guilt. In game-theoretic terms, truthful voting may

be irrational, and the situation that everyone votes truthfully may fail to be a (Bayes-Nash) equilibrium of the voting game.¹⁸ A key goal of this literature is to determine the (often non-majoritarian) voting rules which render truthful voting rational and generate efficient decisions in light of all information spread across voters.

But will voters really engage in such strategic reasoning? The reasoning depends on an extreme motivational assumption: voters care exclusively about the correctness of group outcomes. The rationality of non-truthful voting breaks down quickly as the voter's motivation is enriched by an intrinsic concern for the truthfulness of his own vote, in addition to the concern for the group outcome. For the expressive part of a voter's motivation quickly crowds out the outcome-driven part as the probability of pivotality is usually very small. Besides, even voters with purely outcome-oriented preferences might lack the strategic sophistication for engaging in strategic voting.

9 Concluding assessment

Whether jury theorems are useful for social epistemology and democratic theory is highly controversial. To see clearly through the large variety of proposed theorems, approaches, and objections, we have classified and evaluated theorems based on their premises and their conclusions. Theorems usually make *independence* and *competence* assumptions about voters' correct-voting probabilities; the core question is how these probabilities are understood. Probabilities could be unconditional (Theorem 1), conditional on the state (Theorem 2), or conditional on the full decision problem which includes not just the state but also the common causes of votes (Theorem 3).

Our analysis shows that there is a fundamental tension between independence and competence premises: independence is plausible only for a rich conditionalization, whereas competence is only plausible for a thin conditionalization. Indeed, independence is untenable when construed unconditionally or state-conditionally; but competence is untenable when construed problem-conditionally. As a result, classical jury theorems fail to have jointly justified premises, which is responsible for their (implausible) conclusion that huge groups are infallible. Our suggested response is to use a rich (problem-conditional) conditionalization, while weakening the competence axiom to *tendency to competence*: voters are more often competent than incompetent rather than always competent. The resulting jury theorem no longer concludes that huge groups are infallible; but it still concludes that larger

¹⁸In voting, voters effectively play a strategic game with asymmetric information and common preferences (for objectively correct outcomes).

groups perform better, thereby giving support for majoritarian democracy.

References

- Anderson, Elizabeth. 2006. "The Epistemology of Democracy." *Episteme* 3(1-2): 8–22.
- Austen-Smith, David, and Jeffrey S Banks. 1996. "Information Aggregation, Rationality, and the Condorcet Jury Theorem." *American Political Science Review* 90(1): 34–45.
- Ben-Yashar, Ruth C., and Shmuel I. Nitzan. 1997. "The Optimal Decision Rule for Fixed-Size Committees in Dichotomous Choice Situations: The General Result." *International Economic Review* 38(1): 175.
- Berend, Daniel, and Jacob Paroush. 1998. "When Is Condorcet's Jury Theorem Valid?" *Social Choice and Welfare* 15: 481–88.
- Berend, Daniel, and Luba Sapir. 2005. "Monotonicity in Condorcet Jury Theorem." *Social Choice and Welfare* 24(1): 83–92.
- Black, Duncan. 1958. *The Theory of Committees and Elections*. Cambridge: Cambridge University Press.
- Boland, Philip J, Frank Proschan, and Y L Tong. 1989. "Modelling Dependence in Simple and Indirect Majority Systems." *Journal of Applied Probability* 26(1): 81–88.
- Boland, Philip J. 1989. "Majority Systems and the Condorcet Jury Theorem." *Journal of the Royal Statistical Society. Series D (The Statistician)* 38(3): 181–89.
- Cohen, Joshua. 1986. "An Epistemic Conception of Democracy." *Ethics* 97(1): 26–38.
- Condorcet De, Marquis. 1785. *Essai Sur L'application de L'analyse Á La Probabilité Des Décisions Rendues Á La Pluralité Des Voix*.
- Dietrich, Franz. 2008. "The Premises of Condorcet's Jury Theorem Are Not Simultaneously Justified." *Episteme* 5(1): 56–73.
- Dietrich, Franz. 2006. "General Representation of Epistemically Optimal Procedures." *Social Choice and Welfare* 26(2): 263–83.
- Dietrich, Franz, and Christian List. 2004. "A Model of Jury Decisions Where All Jurors Have the Same Evidence." *Synthese* 142: 175–202.
- Dietrich, Franz, and Kai Spiekermann. 2013a. "Epistemic Democracy with Defensible Premises." *Economics and Philosophy* 29(1): 87–120.
- Dietrich, Franz, and Kai Spiekermann. 2013b. "Independent Opinions?" *Mind* 122(487): 655–85.

- Estlund, David M. 2008. *Democratic Authority: A Philosophical Framework*. Princeton: Princeton University Press.
- Feddersen, Timothy, and Wolfgang Pesendorfer. 1999. "Elections, Information Aggregation, and Strategic Voting." *Proceedings of the National Academy of Sciences* 96(19): 10572–74.
- Goldman, Alvin I. 1999. *Knowledge in a Social World*. Oxford: Clarendon.
- Goodin, Robert E, and Kai Spiekermann. 2012. "Epistemic Aspects of Representative Government." *European Political Science Review* 4(3): 303–25.
- Goodin, Robert E., and Kai Spiekermann. 2015. "Epistemic Solidarity as a Political Strategy." *Episteme* 12(4): 439–57.
- Grofman, Bernard, Guillermo Owen, and Scott L Feld. 1982. "Average Competence, Variability in Individual Competence, and Accuracy of Statistically Pooled Group Decisions." *Psychological Reports* 50: 683–88.
- Grofman, Bernard, Guillermo Owen, and Scott Feld. 1983. "Thirteen Theorems in Search of the Truth." *Theory and Decision* 15: 261–78.
- Kaniovski, Serguei. 2010. "Aggregation of Correlated Votes and Condorcet's Jury Theorem." *Theory and Decision* 69(3): 453–68.
- Ladha, Krishna K. 1992. "The Condorcet Jury Theorem, Free Speech, and Correlated Votes." *American Journal of Political Science* 36(3): 617–34.
- Ladha, Krishna K. 1993. "Condorcet's Jury Theorem in Light of de Finetti's Theorem." *Social Choice and Welfare* 10(1): 69–85.
- Ladha, Krishna K. 1995. "Information Pooling through Majority-Rule Voting: Condorcet's Jury Theorem with Correlated Votes." *Journal of Economic Behavior & Organization* 26(3): 353–72.
- Landemore, Hélène. 2013. *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton, N.J.: Princeton University Press.
- List, Christian, and Robert E. Goodin. 2001. "Epistemic Democracy: Generalizing the Condorcet Jury Theorem." *Journal of Political Philosophy* 9(3): 277–306.
- List, Christian, and Kai Spiekermann. Forthcoming. "The Condorcet Jury Theorem and Voter-Specific Truth." In Goldman and His Critics, eds. Hillary Kornblith and Brian McLaughlin. Wiley.
- Miller, David. 1992. "Deliberative Democracy and Social Choice." *Political Studies* 40(1 Suppl): 54–67.
- Miller, Nicholas R. 1986. "Information, Electorates, and Democracy: Some Extensions and Interpretations of the Condorcet Jury Theorem." In *Information Pooling and Group Decision Making*, eds. Bernard Grofman and Guillermo Owen. Greenwich, Conn.: JAI Press, 173–92.

- Nitzan, Shmuel, and Jacob Paroush. 1982. "Optimal Decision Rules in Uncertain Dichotomous Choice Situations." *International Economic Review* 23(2): 289–97.
- Owen, Guillermo, Bernard Grofman, and Scott L Feld. 1989. "Proving a Distribution-Free Generalization of the Condorcet Jury Theorem." *Mathematical Social Sciences* 17: 1–16.
- Paroush, Jacob. 1998. "Stay Away from Fair Coins : A Condorcet Jury Theorem." *Social Choice and Welfare* 15: 15–20.
- Pivato, Marcus. 2013. "Voting Rules as Statistical Estimators." *Social Choice and Welfare* 40(2): 581–630.
- Pivato, Marcus. 2016. Epistemic Democracy with Correlated Voters. <https://mpra.ub.uni-muenchen.de/69546/> (July 18, 2016).
- Rawls, John. 1971 [rev. ed. 1999]. *A Theory of Justice*. Revised edition. Oxford: Oxford University Press.
- Shapley, Lloyd, and Bernard Grofman. 1984. "Optimizing Group Judgmental Accuracy in the Presence of Interdependencies." *Public Choice* 43(3): 329–43.