politics, philosophy & economics
article

1470-594X
200405 3(2) 131–142

The persuasiveness of democratic majorities

Robert E. Goodin
Australian National University, Australia

David Estlund
Brown University, USA

abstract Under the assumptions of the standard Condorcet Jury Theorem, majority verdicts are virtually certain to be correct if the competence of voters is greater than one-half, and virtually certain to be incorrect if voter competence is less than one-half. But which is the case? Here we turn the Jury Theorem on its head, to provide one way of addressing that question. The same logic implies that, if the outcome saw 60 percent of voters supporting one proposition and 40 percent the other, then average voter competence must either be 0.60 or 0.40. We still have to decide which, but limiting the choice to those two values is a considerable aid in that.

keywords Condorcet Jury Theorem, epistemic democracy, voter competence

The sting in Condorcet's tail

At first blush, the Condorcet Jury Theorem (CJT) seems to be very good news for democracy. That Theorem assures us that, among large electorates, democratic outcomes are virtually certain of tracking the truth, just so long as voters vote independently and are better than random at choosing true propositions, and just so long as there are any 'true propositions' to be found politically.[1]

DOI: 10.1177/1470594X04042960
Robert E. Goodin is joint Professor of Social & Political Theory and of Philosophy in the Research School of Social Science, Australian National University, Canberra ACT 0200, Australia [email: goodinb@coombs.anu.edu.au]
David Estlund is Professor of Philosophy in the Department of Philosophy, Brown University, Providence, RI 02912, USA [email: David_Estlund@brown.edu]

There is, however, a sting in that Theorem's tail: the converse is also true. The same mathematics prove that democratic procedures are virtually certain to yield the wrong results, if voters are less adept than random. That fact has long made democratic theorists wary of embracing epistemic arguments for democracy with the CJT at their heart.[2]

Here we suggest a way of using a previously neglected feature of the CJT as an aid in deciding what to make of democratic majorities, whether we should regard them as persuasive or as the very opposite. We offer no conclusions on that larger question itself. We offer merely a framework for simplifying how to go about addressing it.

The logic of the Jury Theorem

The logic of the CJT is underwritten by the law of large numbers. Given sufficiently large numbers of trials, the law of large numbers tells us that relative frequencies *ex post* will closely approximate probabilities *ex ante*. In the electoral application, given a large number of independent voters, the proportion of people V_i that votes for a proposition ϕ_i will be very near to the probability p_i of each person independently voting for ϕ_i. Thus, if each person is independently 62 percent likely to vote for ϕ_i, then (among a large number of such voters) ϕ_i will win something very close to a 62 percent majority of the votes; if each person is 57 percent likely to vote for ϕ_i, then ϕ_i will win by a 57 percent majority; and so on.

The CJT focuses on ϕ_c, the 'correct, true proposition', and probability p_c (which for convenience is usually assumed to be identical for all voters) that each voter will independently vote for that correct proposition. We know from the law of large numbers that, in any large electorate, V_c will closely approximate p_c. So, in the standard two-option majority-rule case, true proposition ϕ_c is virtually certain of winning ($V_c > 0.5$) just so long as the electorate is sufficiently large and provided the probability of each voter independently voting for the true proposition is $p_c > 0.5$. Conversely, by the same logic, ϕ_c is virtually certain to lose ($V_c < 0.5$) if each voter is independently less likely to be right than wrong ($p_c < 0.5$).

The CJT tells us that the probability that the majority will support the correct option tends toward certainty as the number of voters approaches infinity. The underlying mathematics further reveal that the probability of the majority supporting the correct option is a *rapidly increasing* function of the number of voters, making the CJT strongly applicable to real-sized electorates. Suppose there are two options, and suppose each voter is independently 51 percent likely to choose the correct option (and 49 percent likely to choose the incorrect option): then among a group of 1000 voters, the probability that the majority will vote for the correct option is something very close to 69 percent. If the number of voters is increased to 10,000, then that probability rises to virtual certainty: 99.97 percent.[3] Thus, among electorates of even just moderate-sized towns, much less large nations, the majority is almost certain to choose the right option, just so long as each voter is independently more than half-likely to be right in a two-option choice (but conversely of course if each voter is less than half-likely to be right).

Those results were initially developed, and are most simply stated, for the two-option case. But an analogous result has been proven for the many-option case. Similarly, although the CJT proofs initially assumed voters with identical competence, they have since been extended to voters of varying (but symmetrically distributed) competences. We will say more about both of those issues later. Initially, though, we will discuss these issues in terms of the simpler two-option, identical-competence case.

Assessing voter competence head-on

The CJT seems to be a strong, robust result. The problem is merely in deciding what to make of it — whether to regard it as blessing democratic outcomes or as damning them. In terms of the CJT, that depends purely on whether people on average are better than random, or not, in making political decisions involving matters of fact.

It is difficult to address that question head-on, because there is much to be said on both sides of that question:

- Most of the propositions that are put to a vote do not admit of any easy, objective test of truth. The fact that there is a good case to be made on both sides is precisely why we put the issue to a vote in the first place.
- Survey research shows that the electorate is woefully ill-informed about most matters of public affairs. How much to make of that is unclear, however. There may be mechanisms of 'low-information rationality' that can guide even ill-informed voters to rational choices.[4] Besides, some might say, it is hard to imagine how people could do systematically worse than random at picking the right answers, even if they were totally uninformed.
- If, however, voters all cued on the same sources of information, and those sources were misleading, then there might be a 'common mode failure' that could well lead voters individually and hence collectively to perform systematically worse than random.

In short, the question of whether voters are on average better than random is not easily decidable when approached head-on.

Inferring voter competencies from electoral outcomes

Our suggestion is to go roundabout. Instead of trying to decide whether or not voters are better than random, in general, let us use the outcome of the election to help us reflect on judgements about average voter competence, and hence on the decision of whether to place positive or negative faith in the outcome of the democratic election.

The CJT uses the law of large numbers to pass from assumptions about p_c (voter competence) to conclusions about V_c (the share of the vote won by the correct outcome). Here we propose to work in the reverse direction, deriving inferences about voter competence (p_c) from the actual distributions of votes (V_i).

The law of large numbers works in both directions. Given a value for p_c and a large electorate, it allows us to predict the value of V_c. That is the key to the CJT story as it is ordinarily told. By the same token, however, once an election has been held and we know the distribution of votes, V_i, we can infer from the distribution of values of V_i two alternate possible values for p_c.

The law of large numbers says that, among a large number of independent voters each equally likely to vote for the *i*th proposition ϕ_i, the probability of each voting for ϕ_i is the same as the proportion of them who did vote for ϕ_i, or $p_i = V_i$.

Suppose, for convenience, that each voter has the same probability as each other of voting for the correct outcome, and each votes independently of every other. And suppose, again purely for expository convenience, that we are dealing with a standard two-option majority-rule situation. (We discuss applications to *k*-option cases below.)

What we can infer from the distribution of a large number of votes is a *pair* of possible values of voter competence, p_c. One of two possible conclusions about voter competence we can infer from a distribution of V_i concerns the case where voters are more likely to vote for the correct option than any other. In the two-option case, this amounts to saying $p_c > 0.5$. From the law of large numbers, we know that if $p_c > 0.5$ then the correct outcome is virtually certain to win. Also from the law of large numbers, we know that the probability of voters voting for the correct outcome (p_c) is, in this first case, the same as the proportion of voters voting for that winning outcome w, V_w. Thus, one possible value of p_c is V_w.

The other possible value of p_c is associated with the case where voters are more likely to be wrong than right. In the two-option case, that amounts to saying that $p_c < 0.5$. By the law of large numbers we know that, in the two-option case, if $p_c < 0.5$ then the correct outcome is virtually certain to lose. Furthermore, the law of large numbers allows us to infer the probability of voters voting for the correct outcome (which is in this case the losing outcome). That is the same as the proportion of voters voting for the losing outcome, which in the two-option case is $(1 - V_w)$. The second possible value of p_c is, thus, $p_c = 1 - V_w$.

Suppose for example 60 percent of a large electorate votes for one option and 40 percent for the other. From that distribution of votes, together with standard CJT assumptions about uniform competence and independent voters, we can infer that the probability of each voter voting for the correct outcome (p_c) must be *either* around 60 percent *or* around 40 percent.

That simplifies things nicely, when it comes to trying to decide what to make of the outcome of that election — whether to conclude that the majority verdict is almost certainly right or almost certainly wrong. We have merely to decide which of those two alternative possible values of voter competence we believe to be more likely, in the case at hand. Is it more credible that, on this subject, voters on average are 60 percent likely to be right? Or is it more credible that voters on average are only 40 percent likely to be right?

Different competences on different issues

The CJT typically works with simplifying assumptions of uniform voter competence in two dimensions. First, it is typically assumed that each voter is as competent as each other voter; and second, these discussions typically proceed as if voters were equally competent across all subjects.

Neither assumption is strictly necessary. The first assumption can be relaxed, as we have already said. Voters need not have identical competence; the CJT can be re-proven using the mean competence of voters of varying competence, just so long as the distribution of voter competences is symmetrical around the mean.[5]

The second assumption is formally even more dispensable, of course. The CJT's basic conclusion — that democracy is a good truth-tracker — holds just so long as the mean voter is 'better than random' on each topic. There is no need for them to be equally competent across all topics.

Nevertheless, here is a clear and interesting implication of our way of using the law of large numbers in the reverse direction to ordinary CJT derivations. From the simple fact that different elections are decided by different margins, it immediately follows that voter competences *must* differ across different subject matters (insofar as voters vote purely on the basis of perceived truths).

Letting landslides matter

In the standard CJT, we do not need to know exactly what average voter competence actually is. All we really need to know is whether it is greater than or less than 0.5. In the former case, the majority is almost certain to be correct, in a sufficiently large electorate.[6] In the latter case, it is almost certain to be incorrect.

That is certainly a strong result. There is, however, a certain embarrassment surrounding it. The embarrassment is that that result is utterly impervious to the *proportion* of the electorate voting for or against the proposition. Given a large number of electors (which the CJT requires to work at all), there is no material difference between a 51 percent victory and a 70 percent victory. Given a very large number of voters, the first is (very, very nearly) as certain to be correct as the latter. That seems an odd conclusion.

The reason for that conclusion is that what drives the CJT is the absolute majority in favour of one option rather than the other, not the relative majority. If there are 10,000 more voters for proposition ϕ_1 than ϕ_2 then the CJT question is, in effect, 'How can 10,000 people be wrong?' And on standard CJT models, the question remains equally pressing whether it is 10,000 voters out of 100,000 or 10,000 out of 100,000,000,000. That is just the way that the mathematics of the CJT work.[7]

Working within the standard CJT framework, the only way to avoid the conclusion that ϕ_1 (with its extra 10,000 supporters) is virtually certainly correct is to say that $p_c < 0.5$ — in which case the CJT insists with equal confidence that not-ϕ_1 (which is to say, ϕ_2) is virtually certain to be correct instead. There is simply no scope, within standard CJT models applied to very large electorates, for regarding a 51 to 49 percent outcome as appreciably less conclusive than a 70 to 30 outcome.

The standard CJT model presses us to decide in advance of the vote (or anyway independently of the vote) whether voters on average are more or less likely to be right than random. Everything else follows from that, together with the law of large numbers. Our alternative approach is to start with the outcome of the vote and work back (via the law of large numbers) to a pointed choice between alternative hypotheses about average voter competence on the sort of issue at hand.

Whereas the CJT forces us to accept either one strong conclusion (the majority is almost certainly right) or another (the majority is almost certainly wrong), our approach allows a more balanced assessment, taking account of the proportion of the electorate voting each way. Suppose that we know the electorate has split 70 percent to 30 percent on this issue. We know that the electorate is large enough for those figures to track the probability that each (or the average) voter has of being right on this issue. We just have to decide whether the average voter is more likely to be 70 percent reliable on this subject or 30 percent.

Some might suppose that, as a sort of extension and generalization of the 'principle of charity', we ought (unless we have some special story to tell about how this case is unusual) ordinarily suppose that our fellow citizens are more likely to be right than wrong, and therefore that the winning outcome is quite probably the correct one.[8]

Those who charitably trust the veracity of others' reports will ordinarily be a lot more confident in a proposition where the proportion of equally trustworthy voters siding with the majority is large rather than small. If the vote is 50.1 percent to 49.9 percent, and they have to choose between *that* pair of possible voter competences, even the charitably inclined must be less confident that the winning proposition is the right one. The charitable may regard it as incredible (absent some special reasons for being suspicious of the majority, more of which below) that average voter competence is only 0.30, in the 70 : 30 breakdown of the vote. But when it comes to interpreting the 49.9 : 50.1 percentage breakdown of the vote, even the charitable may suppose it is not all that incredible that people on average are more likely to be 0.499 as opposed to 0.501 reliable.

Of course, either way you run the story — the standard CJT way, or ours — the fact remains that, in a large electorate, a majority of better-than-random truth-trackers is virtually certain to be correct in their judgement. If voters are $p_c > 0.5$, no matter by how much or little p_c exceeds 0.5, then in any very large electorate the correct outcome is very, very likely to win.

Our way of running the story, however, might give those inclined toward a principle of charity some principled rationale for attaching varying degrees of confidence to the proposition that voters are indeed, on average, better-than-random truth-trackers. The issue we pose is, what 'degree of confidence' do we have in the proposition, '$p_c > 0.5$'? Those inclined toward a principle of charity would respond by saying that, *ceteris paribus*, we ought ordinarily have more confidence in that proposition where the choice of alternative possible voter competences is between a very high and a very low one. They would respond that they have much less confidence in that proposition where the choice is between two alternative values for p_c that are very close to one another. And remember, those alternative possible values of p_c are derived, on our approach, from the proportions of people actually voting for each option. So that amounts to saying that, *ceteris paribus*, larger proportional majorities will seem more persuasive for those inclined to apply a principle of charity in the case at hand.

Being suspicious of majorities

Whether or not we ought to apply a principle of charity, either in general or especially in any particular case, is however an open question. Sometimes we might have good grounds for not doing so, that is to say, for being suspicious of the majority (and all the more suspicious of larger majorities than smaller ones).

Majorities might reflect nothing more than systematic biases within the electorate, and bias ought epistemically be discounted. Our approach provides some scope for detecting bias, and for discounting the results of majority votes accordingly. Bias would show up in these analytics as 'bloc voting': a failure of the assumption, crucial to all CJT-style results, that every voter's vote be statistically independent of every other's.

As we have already remarked, different elections are won by different majorities. On our approach, that implies that voter competences vary from issue to issue (insofar as voters are trying to track the truth at all when voting on those issues). Sometimes, however, inspecting the distribution of the vote will suggest an alternative explanation than that voters are more (or less) competent than usual on this issue.

Suppose we are dealing with some racially charged issue; suppose that 90 percent of voters back option ϕ_1 and 10 percent back option ϕ_2 (i.e. not-ϕ_1); and suppose demographic breakdown of the electorate is 90 percent white and 10 percent black. One's suspicions would certainly be aroused. One would certainly require a lot of reassurance (by looking at district-by-district breakdowns of the vote, etc.) that voters were 90 percent accurate at tracking the truth, rather than just 100 percent accurate at tracking race.

There are of course plenty of other ways for one's suspicions to be aroused in these and cognate cases. But looking at the proportional breakdown of the vote, and asking ourselves whether that is more plausibly a reflection of voter competence or voter interests and biases, is a useful supplement to the ordinary battery of bias-detectors.

Our approach also provides similar scope for detecting breakdowns in the 'independence' assumption so crucial to CJT derivations. Suppose the vote is 98 to 2, as it was in the US Senate vote on the Gulf of Tonkin resolution. It hardly seems credible that, on virtually any subject that comes before politicians for a vote, people are $p_c > 0.98$ likely to be right. It seems just as incredible to suppose that they are only $p_c > 0.02$ likely to be right. The most plausible interpretation, in a case like that, is that the voters were not judging the matter independently of one another but that they were, instead, all basing their votes on the same biased source, President Johnson's fabrications.

Beyond two options

So far we have, for simplicity, been following the CJT tradition of talking in terms of majority voting over two options. But, as has been proven, the CJT can be extended to plurality voting over multiple options.[9] In the two-option case, each voter has to be $p_c > 0.5$ likely to vote for the right outcome, in order for the law of large numbers to assure us that the outcome winning the majority is virtually certain to be right. The corresponding requirement in the k-option case is that each voter be more likely to vote for the correct option than any other option.[10]

Whereas the standard CJT analysis, in that case as in the other, infers the likelihood of democratic pluralities being correct from assumptions about voter competence, we once again would infer alternative voter competences from the distribution of the votes.

Here, as before, there are two cases to consider. One is that on average voters are more likely to vote for the right option than any other among the k-options, and hence the plurality winner is the correct outcome. In that case, from the law of large numbers, average voter competence is once again just equal to the proportion of the votes the plurality winner secured ($p_c = V_w$).

The second case to consider is that on average voters are more likely to vote for a wrong option than the right one. Then one of the (k-1) defeated options is actually the right outcome, and average voter competence is equal to the proportion of the votes that that option secured. The difficulty, in the k-option case, lies in knowing *which* of the (k-1) defeated options is the correct one. All we can know in the k-option case is that the alternative voter competence p_c associated with the possibility that voters are more likely wrong than right corresponds to the vote share V_i of one or other of the (k-1) defeated options.

For practical purposes, perhaps the most useful way of framing that is to say that we are faced with the following choice: *either* we can conclude that average voter competence equals the vote share of the plurality winner; *or else* we can conclude that average voter competence is no greater than the vote share of the most popular of the defeated alternatives. Thus, in a 4-option election with a large number of voters, with a distribution of the votes in the ratios 40 : 30 : 15 : 15, we can say that average voter competence must either be 40 percent or not greater than 30 percent. It might be as low as 15 percent, of course. But just forcing us to think of it as 'not greater than 30 percent' might help focus our thinking.

Reconceptualizing voter competence over many options

Where votes are split among many options, the option that wins the plurality of such votes might nonetheless attract a relatively small proportion of the total vote. The CJT tells us with almost complete confidence that it is the right outcome nonetheless, just so long as each voter is more likely to vote for the correct option than any other.

The implication for inferences about voter competence might nonetheless seem awkward. Imagine an 11-option case, in which the vote shares won were 12 percent for one option and (just to make things simple) 8 percent for each of the other options. Then our approach would lead us to conclude that average voter competence is either 12 percent or 8 percent. Neither seems very impressively high.

Of course, one way to console ourselves would be to recall (as we have already remarked) that voter competences must be contextualized to issues and choice situations. And if voters were 12 percent likely to choose the correct option in that particular 11-option case, and no more likely to choose any other option, then that is enough to guarantee in CJT fashion that a large number of such voters are almost certainly right in their plurality verdict. So contextualizing to cases is one way to make the result seem more impressive, despite $p_c = 0.12$ seeming to be such a low number.

Here is another way of looking at that situation. Suppose we were conducting a run-off style election, with the option winning fewest votes being eliminated in each successive round. (And in cases of a tie, as in our 11-option case above, which of the equally unpopular options is eliminated is then decided randomly.) Let us continue to assume, for convenience, that voters are all equally competent. So the voters who in the first round had voted for the option which has now been eliminated in the second round will distribute themselves among the correct option and each of the remaining nine incorrect options in the same proportions; and so on, as low-vote options are successively eliminated in subsequent rounds of the run-off. By the time we get down to the final round of this run-off, voters will be distributed 60 : 40 between the two remaining options.

In that way, 12 : 8 in the 11-option case we described can be regarded as equivalent to 60 : 40 in the sort of 2-option case with which we are more familiar.

Thinking of it that way, the numbers are much more impressive, in both respects. First, saying that average voter competence is either $p_c = 0.60$ or $p_c = 0.40$ is to say that it might be really pretty high ($p_c = 0.60$ certainly *looks* a lot better than $p_c = 0.12$, even though as we have shown it can be regarded as its 2-option equivalent). Second, to say that average voter competence is either $p_c = 0.60$ or $p_c = 0.40$ is to say it is either pretty high or pretty low, making it a starker choice than it seemed at first brush when we began thinking of it as either $p_c = 0.12$ or $p_c = 0.08$.

Conclusion

How much in the end has all this really simplified the problem of deciding what to make, epistemically, of a democratic verdict? After all, in the standard CJT we only have to decide whether or not voters on average are more likely to be right than random. If they are, and the other assumptions of the CJT are met, then the verdict of a large electorate is epistemically compelling; if they are not, then their verdict is almost certain to be epistemically in error. To decide which interpretation is correct, in the standard CJT case, we need only decide which side of 'random' the average voter is on.

Of course, average voter competence may well be right around 'random'. Some of the considerations suggest that it might be a little better than random. Others suggest that it might be a little worse than random. Where exactly the balance falls is an open question. But that, of course, is the *key* question we have to resolve in deciding what to make, epistemically, of the democratic verdict.

In our reframing of the CJT results, the question can sometimes be posed in a more stark way. Knowing that the democratic outcome was 60:40, we have then to decide which possibility is more credible. Is it more credible that in this sort of case the average voter is 60 percent likely to choose correctly (in which case the verdict of the majority is almost certain to be epistemically correct)? Or is it more credible that the average voter is in this sort of case only 40 percent likely to choose correctly (in which case the majority verdict is almost certain to be epistemically incorrect)?

In the case of close elections (51:49, for example) the same problem as with the standard CJT re-emerges, of course. There we have to decide whether it is more credible that the average voter is 51 percent likely to choose correctly, or whether it is more credible that average voter competence is 49 percent. For reasons given at the outset, that is a hard one to call. And because we are uncertain on that issue, we are uncertain also just how much epistemic credence to place in the verdict of such a slim majority, even if the size of the electorate is very large.

But that seems right. Narrow electoral victories *should* pack less of a punch, epistemically and perhaps democratically as well, than should massive majorities. It is an unpleasantly counter-intuitive feature of the CJT, as ordinarily construed, that a 51:49 majority is very, very nearly as likely to be correct as is the majority in a 70:30 landslide, just so long as the electorate is large. It is one of the great benefits of our approach that it avoids that counter-intuitive result and leaves room for being more skeptical of slim proportional majorities than massive ones.

notes

We are grateful to Dan Hausman for asking the question that led us down this line of enquiry and for subsequent comments from Geoff Brennan, John Dryzek, Frank Jackson, Christian List and Gerry Mackie. This article was written during Estlund's year as a Harsanyi Fellow in the Social & Political Theory Program, Research School of Social Sciences, Australian National University.

1. Formally, the probability that the electorate collectively will reach a correct decision is:

$$P_N = \sum_{i=m}^{N} (N!/[N-i]!i!)\,(p_c)^i\,(1-p_c)^{N-i}$$

   where N is the number of voters, m is the number of votes required to win and p_c is each voter's own probability of being correct (which for convenience is assumed to be identical for all voters, whose votes are also assumed to be statistically independent of each other's).
2. Jeremy Waldron, 'Democratic Theory and the Public Interest: Condorcet and Rousseau Revisited', *American Political Science Review* 83 (1988): 1322–1323 and *Law and Disagreement* (Oxford: Clarendon Press, 1999), p. 135. David Estlund, 'Making Truth Safe for Democracy', in *The Idea of Democracy*, edited by David Copp, Jean Hampton and John E. Roemer (New York: Cambridge University Press, 1993): 71–100, p. 93. Gerald Gaus, 'Does Democracy Reveal the Voice of the People? Four Takes on Rousseau', *Australasian Journal of Philosophy* 75 (1997): 141–162, p. 150 and *Contemporary Theories of Liberalism* (London: Sage, 2003), pp. 158–164.
3. Brian Barry, 'The Public Interest', *Proceedings of the Aristotelian Society* (*Supplement*) 38 (1964): 1–18, p. 9, drawing on Duncan Black, *The Theory of Committees and Elections* (Cambridge: Cambridge University Press, 1958), pp. 163–165.
4. Arthur W. Lupia and Matthew D. McCubbins, *The Democratic Dilemma: Can Citizens Learn What They Need to Know?* (Cambridge: Cambridge University Press, 1998).
5. Bernard Grofman, Guillermo Owen and Scott L. Feld, 'Thirteen Theorems in Search of the Truth', *Theory & Decision* 15 (1983): 261–278. Given that the distribution of p is both left- and right-censored ($0 \leq p \leq 1$), the distribution can be literally symmetrical for cases of mean competence other than 0.5 only if there are literally no voter competences further from the mean on the 'long' side than the distance to the censor (0 or 1) on the 'short' side. Thus, e.g., if mean competence is

0.7, then the distribution can be literally symmetrical only if there is no voter with competence below 0.4. Hence the assumption of symmetry here might be more demanding than it first appears.

6. What is 'sufficiently large' depends on what probability you want to treat as 'almost certain' and on how much above 0.5 average individual voter competence is.
7. The same is true of Bayesian analogues to the CJT. See Robert E. Goodin, 'The Paradox of Persisting Opposition', *Politics, Philosophy & Economics* 1 (2002): 109–146 and *Reflective Democracy* (Oxford: Oxford University Press, 2003), Ch. 6 and, more formally, Christian List, 'On the Significance of the Absolute Margin', *British Journal for the Philosophy of Science*, forthcoming.
8. The 'principle of charity' is explicitly appealed to as grounds for trusting the testimony of others, in different fashions, by both C.A.J. Coady, *Testimony* (Oxford: Clarendon Press, 1992), Ch. 9 and Tyler Burge, 'Content Preservation', *Philosophical Review* 102 (1993): 457–488, p. 487. Note that, however similar in spirit, this cannot be literally the same 'principle of charity' as Davidson's, as shown by Gary Ebbs, 'Learning from Others', *Nous* 36 (2002): 525–549.
9. Christian List and Robert E. Goodin, 'Epistemic Democracy: Generalizing the Condorcet Jury Theorem', *Journal of Political Philosophy* 9 (2001): 276–306. Goodin, *Reflective Democracy*, Ch. 5.
10. In the *k*-option case it is necessary (in a way that was otiose to add in the 2-option case) that each voter is more likely to vote for the correct outcome *than any other*, rather than just saying $p_c > (1/k)$. This stipulation is needed to protect against this sort of scenario: imagine a choice among four options; people are $p_1 = 0.3$ likely to vote for option 1 (the correct option), thus satisfying the $p_1 > (1/k)$ requirement; but people are also $p_2 = 0.4$ likely to vote for option 2 (one of the incorrect options) and $p_3 = p_4 = 0.15$ likely to vote for each of the other two incorrect options (options 3 and 4, respectively). The law of large numbers tells us that the most popular wrong option (option 2) will win the plurality vote over the correct option, by a plurality of 40 percent to 30 percent of the vote. Hence the need to add the further stipulation in the *k*-option case.