

A (mainly epistemic) case for multiple-vote majority rule

Richard Bradley and Christopher Thompson
Department of Philosophy, Logic and Scientific Method
London School of Economics
Houghton Street
London WC2A 2AE
r.bradley@lse.ac.uk

January 6, 2011

Abstract

Multiple-vote majority rule is a procedure for making group decisions in which individuals weight their votes on issues in accordance with how competent they are on them. When individuals are motivated by the truth and know their relative competence on different issues, multiple-vote majority rule performs nearly as well, epistemically speaking, as rule by an expert oligarchy, but is still acceptable from the point of view of equal participation in the political process.

1 Introduction

A voting rule is a system for producing group decisions or judgements on propositions or proposals as a function of how the individuals making up the group vote on these propositions or proposals. Different voting systems display quite different virtues: epistemic ones, such as the ability to produce the correct decision or to avoid very bad ones; procedural ones such as ensuring equal participation in decision making, safeguarding individual rights and recognising individual autonomy; and practical ones, such as flexibility or ease of implementation. The paper will focus on just two of these desiderata: the propensity of the rule to produce correct decisions and the extent to which it ensures equality of participation. Not because they exhaust the relevant characteristics of a voting rule, but because they enable us to display in a precise way the trade-off between the goal of involving in decisions as many as possible of the people that are affected by them and the goal of reliable and prompt decision-making.

The trade-off is somewhat obscured by over-optimistic interpretations of the Condorcet Jury Theorem that, by focusing on cases in which voters are large in number and at least minimally competent, suggest that democratic forms of decision making deliver the best of both worlds. The core of the epistemic defence of democracy is the claim that in giving all an equal say, majority voting delivers group decisions that reliably track the truth; indeed more reliably than dictatorial or oligarchic ones. However in cases in which voters are modest in number, and the competence of individual voters varies to a large degree from voter to voter and issue to issue, voting systems that

give unequal weight to those that are more competent will produce better decisions, an epistemic benefit that is purchased at the cost of unequal participation.

In this paper we evaluate a rule that we dub multiple-vote majority rule (MVMR) and argue that, under suitable considerations, it achieves a better balance of epistemic reliability and equality of participation than several other well-known rules, including majority rule. The rule of multiple-vote majority rule works as follows. Suppose that the group must make a decision on a number of propositions, say ten in all. Then each individual is allocated ten votes which they can place on any of the propositions. That is, they can vote up to ten times on a single proposition as long as the total number of votes they cast does not exceed ten (i.e. the individual abstains from voting on the remaining nine propositions). The collective decision is reached by adding up the number of votes for and against the proposition and accepting it if and only if the number for exceeds the number against. If individuals are only interested in the truth, they will try and place their votes where they will do the most good i.e. will have the greatest positive impact on the group competence. For instance, if they consider themselves more competent on two of the propositions than on the others, they could concentrate their votes on them. Our conjecture is that when competence varies both between individuals and between propositions, and individuals have second-order competence on the question as to what they are competent on, then multiple-vote majority rule can produce greater group competence than majority rule without sacrificing equality of participation.

Our discussion works within a framework inherited from well-known epistemic defences of democracy (see for instance Cohen (1986), Estlund (1997) and List and Goodin (2001)). Its essential constituents are as follows. A group of individuals hold diverse judgements on one or more propositions, but must reach a collective decision as to whether to accept or reject each one of them. Each proposition can be true or false and each individual's judgement on each proposition is assumed to be reliable to some degree, this being measured by the conditional probability that they will vote for the proposition, given that it is true. Each individual's acceptance of the proposition is signalled by a vote for the proposition, rejection by a vote against.

A voting rule is a proposition-independent procedure for determining collective acceptance or rejection of any proposition on the basis of how the individuals vote and any relevant characteristics of the individual voters. The only characteristic we consider is the individuals' competence on the propositions on which they vote. A rule that ignores these characteristics and determines the social acceptance or rejection of the propositions solely on the basis of the pattern of votes will be called an anonymous rule. Anonymity is desirable from a procedural point of view because it encodes equality of participation. On the other hand, what we are interested in, from an epistemic point of view, is the capacity of voting rules to arrive at the correct verdict i.e. to accept propositions when they are true and reject them when they are not. Whether a rule is good at doing this depends, of course, on how competent individuals are, but also on how the rule exploits this competence.

More formally, we define the reliability of a particular voting rule to be a function from a profile of individuals' reliabilities to the corresponding group reliability, this being the conditional probability that the rule in question will lead to acceptance of the proposition, given that it is true. The tendency of a rule to produce the correct outcome may then be defined as the expected value of its reliability, relative to some probability measure over profiles of individual reliabilities.

To study the reliability of various voting rules, we look at three cases. First we look at a baseline case in which individuals are assumed to have the same reliability as each other and on every proposition. This case has been already been studied elsewhere

(by List (2008) for instance) and will simply rehearse existing results. In the second treatment, we drop the assumption of homogeneous reliability amongst individuals and allow that the reliability of individuals is heterogeneous. In the third treatment, we also drop the assumption of homogeneous reliability across propositions.

Throughout we will make a number of important assumptions. Firstly, we assume that individuals are only interested in the truth and so they vote for a proposition only if they believe it to be true, or more likely than not to be true. When individuals vote for other reasons - for instance, because a proposition being collectively accepted benefits them in some way - then many of the conclusions we draw here will have to be modified. Multiple-vote majority rule remains a rule of interest in these circumstances because it offers a way of allowing individuals to express preference intensities by placing more votes on propositions that they feel strongly about. But we put aside this discussion here.

Secondly, we assume that individuals' votes are independent of one another in the sense that the probability that any voter will accept or reject any proposition is independent of whether other voters do. Although unrealistic, this assumption is often made in the literature on epistemic democracy and we adopt it here partly for reasons of simplicity and partly to facilitate comparisons with existing results. We do not think that the main conclusions of the paper would be substantially affected by dropping it.

Thirdly, we assume that individuals' votes on each of propositions in the given set are independent in the sense that the probability that an individual will accept any proposition in the set is independent of whether she accepts or rejects any other. This assumption will not generally be satisfied in real voting situations as the logical independence of propositions does not by any means ensure their probabilistic independence. We adopt it for the same reason as the previous independence assumption, namely simplicity and ease of comparison with existing results. And in this case too we do not think that the main conclusions of the paper will be much changed by a more realistic treatment.

Finally, we assume that voters know their own relative reliability on propositions, but have little knowledge of the reliability of others. Whether this assumption holds is, of course, a contingent matter. There is some evidence of a correlation between actual competence and self-perception of competence in a number of domains¹, but it's clearly unrealistic to suppose that in general individuals have very precise knowledge of this kind. But although we use precise values for relative reliability in the model of voting with heterogeneous reliability (see appendix), the argument does not require that agents' know these precise values. What it does require is that they can make qualitative comparisons between their reliability on different propositions, so that they can shift votes away from the issues on which they are less reliable and towards those on which they are more confident of their judgemental abilities.

¹See for instance Kruger and Dunning (2009). They show that for issues of logic and humor (though not grammar), while the top percentiles tended to under-estimate their ability and the bottom percentiles tended to over-estimate their abilities, there is nevertheless a strong relationship between the actual ability of agents and their perceived ability. Extrapolating from this result, if an agent is more competent on some questions than on others, they should be able to detect this difference. As a consequence they should be able to allocate more of their votes to the questions they feel most competent on, at the expense of those questions they feel the least competent on.

2 Homogeneous Reliability

Suppose that a group of n individuals must decide on whether each member of a set of five propositions, $J = \{P_1, \dots, P_5\}$, is true or not. Let the positive reliability r_j^i of individual i on the j th proposition be defined as the probability that i will vote for proposition P_j , given that the proposition is true. The negative reliability \bar{r}_j^i of i on P_j is the probability that i will vote against P_j , given that P_j is false. The competence of i on P_j , denoted by c_j^i , is measured by the ratio of the probability that they will vote for P_j , given that it is true, to the probability that they will vote for it, given that it is false, i.e.:

$$c_j^i = \frac{r_j^i}{1 - \bar{r}_j^i}$$

In virtue of the assumption that individuals' votes on the propositions are probabilistically independent, the positive reliability of individual i on the whole set of propositions (or equivalently their conjunction), r_0^i , must equal $\prod_{j=1}^5 r_j^i$. (We won't study negative reliabilities on sets of propositions, but it is simple enough to see how the treatment of positive reliabilities extends to them).

The group's reliability respectively on proposition j and on the whole set of propositions will be denoted by r_j^G and r_0^G . These group reliabilities depend on the voting rule employed to determine acceptance of any of the propositions. Group acceptance of the set of propositions can however be determined in two ways: by applying the voting rule to each proposition in the set and then accepting the whole set only if every one of its members is accepted (the premise-based procedure) or by applying the voting procedure to the set itself, or more exactly, the conjunction of its members (the conclusion-based procedure). These two procedures will not generally produce the same result, as we show below.

In our first treatment we assume homogeneous competence in the group and on all five propositions in virtue of identical positive and negative reliabilities. Let the uniform positive reliability of individuals be r , so that the individuals' common reliability on the set of propositions equals r^5 , and that of the group be r^G . A number of different voting rules can now be compared in terms the values of r_j^G and r_0^G that they yield, as well as the probability of the group reaching a correct verdict on all 5 propositions.

1. *Dictatorship*: On this procedure the group accepts a proposition iff some individual - the dictator - votes for it. Group reliability in this case just equals that of the dictator. Given the assumption of homogeneous competence on propositions this means that $r_j^G = r$ and $r_0^G = r^5$.
2. *Unanimity*: On this rule, the group accepts a proposition j if and only if, every individual in the group votes for j . Given the assumption of independence, group reliability is the product of individual competencies. Hence with homogeneous reliability:

$$\begin{aligned} r_j^G &= r^n \\ r_0^G &= r^{5n} \end{aligned}$$

3. *Majority Rule*: The group accepts a proposition if and only if a majority of individuals in the group vote for it. Given independence and homogeneous reliability this means that:

$$r_j^G = \sum_{h > \frac{n}{2}} \binom{n}{h} r^h (1-r)^{n-h} \quad (1)$$

If a premise-based procedure is employed to determine group acceptance or rejection of the set of propositions, then independence and uniform reliability on propositions implies that:

$$r_0^G = (r_j^G)^5$$

On the other hand, under a conclusion-based procedure, these assumptions imply that²:

$$r_0^G = \sum_{h > \frac{n}{2}} \binom{n}{h} r^{5h} (1 - r^5)^{(n-h)} \quad (2)$$

How then do these rules compare in terms of the group reliability they induce? The table below displays the probabilities, for each of the rules, of delivering the correct verdict on a single proposition (SP) and on the whole set of them (MP), for a range of values for positive reliability r and population size n . In the case of majority rule, the table gives group reliabilities under both the premise- and conclusion-based procedures (MP(P) and MP(C) respectively). All values in tables 1-6 in the following are rounded.

Table 1: Homogeneous Reliability, $r = 0.49$

n		Dictator	Unanimity	Majority rule
5	<i>SP</i>	0.49	2.82×10^{-2}	0.48
	<i>MP(P)</i>	2.82×10^{-2}	1.80×10^{-8}	2.58×10^{-2}
	<i>MP(C)</i>	–	–	2.16×10^{-4}
15	<i>SP</i>	0.49	2.25×10^{-5}	0.47
	<i>MP(P)</i>	2.82×10^{-2}	5.82×10^{-24}	2.29×10^{-2}
	<i>MP(C)</i>	–	–	0
105	<i>SP</i>	0.49	2.96×10^{-33}	0.42
	<i>MP(P)</i>	2.82×10^{-2}	2.25×10^{-163}	1.31×10^{-2}
	<i>MP(C)</i>	–	–	0
1005	<i>SP</i>	0.49	4.44×10^{-312}	0.26
	<i>MP(P)</i>	2.82×10^{-2}	1.72×10^{-1557}	1.18×10^{-3}
	<i>MP(C)</i>	–	–	0
10005	<i>SP</i>	0.49	2.58×10^{-3100}	0.02
	<i>MP(P)</i>	2.82×10^{-2}	1.15×10^{-15498}	6.46×10^{-9}
	<i>MP(C)</i>	–	–	0

²The reliability of an individual on the conclusion (which is a conjunction of different premises) will be no higher, and typically much lower, than her reliability on any of the premises. For convenience we assume that the reliability of the conjunction of h independent propositions is $(r_j^i)^h$.

Table 2: Homogeneous Reliability, $r = 0.51$

n		Dictator	Unanimity	Majority rule
5	SP	0.51	3.45×10^{-2}	0.52
	$MP(P)$	3.45×10^{-2}	4.89×10^{-8}	3.76×10^{-2}
	$MP(C)$	–	–	3.90×10^{-4}
15	SP	0.51	4.12×10^{-5}	0.53
	$MP(P)$	3.45×10^{-2}	1.17×10^{-22}	4.21×10^{-2}
	$MP(C)$	–	–	0
105	SP	0.51	1.97×10^{-31}	0.58
	$MP(P)$	3.45×10^{-2}	2.98×10^{-154}	6.56×10^{-2}
	$MP(C)$	–	–	0
1005	SP	0.51	1.28×10^{-294}	0.74
	$MP(P)$	3.45×10^{-2}	3.45×10^{-1470}	0.22
	$MP(C)$	–	–	0
10005	SP	0.51	1.74×10^{-2926}	0.98
	$MP(P)$	3.45×10^{-2}	1.58×10^{-14629}	0.89
	$MP(C)$	–	–	0

No voting rule performs better epistemically than all others for every level of individual reliability. Consider first the single proposition case. If $r > 0.5$, then the Unanimity rule is beaten by Dictatorship which is beaten by Majority rule (once adjustments for ties have been made). If $r < 0.5$, then the Unanimity rule is beaten by Majority rule which is beaten by Dictatorship. The reason for this is the law of large numbers, the mechanism underlying the Condorcet Jury Theorem: as the population size rises, the probability that the proportion of individuals voting for a true proposition P will equal the population's reliability on P tends to one. Hence if r is greater than one-half, majority rule will accept true propositions with high probability. On the other hand, if r is less than one-half, then majority rule will lead with very high probability to rejection of true propositions.

Defenders of majority rule claim that the former case is to be expected as even poorly informed individuals are likely to be better than purely random devices in picking true propositions. But this argument cannot be right in general, as the case of multiple propositions shows. If individuals have a reliability of, say, 0.51 on each of the five propositions, then their reliability on the conjunction of them is $0.51^5 = 3.4503 \times 10^{-2}$. Hence, in accordance with the predictions of the Condorcet Jury Theorem, as the number of voters increases the probability of a correct majority verdict tends towards zero. This explains the poor epistemic performance of conclusion-based majority rule. For this reason, we dispense with conclusion-based voting for the remainder of the paper and focus on premise-based voting, whereby the conclusion is accepted by the group iff each of premises is accepted by a majority.

In both cases of reliability above and below a half the Unanimity rule will accept true propositions with only extremely low probability. (On the other hand, it will also accept false propositions with low probability, so Unanimity will be a good rule when the consequences of accepting a false proposition are disastrous).

3 Heterogeneous Reliability

In this section we drop the assumption that the reliabilities of individuals are the same and allow for heterogeneity in the population. Voting rules will again be compared in terms of the group reliability they induce both on single propositions and on a set of them. In this richer setting moreover we can meaningfully look at new rules which exploit reliability heterogeneity by putting decisions in the hands of the more or most reliable.

There are two salient categories of such rule: ‘oligarchic’ rules which use only the votes of the most reliable and ‘weighting’ rules which weight votes according to reliability and accept a proposition only if the weighted sum of votes for it exceeds some threshold. The former category includes expert dictatorship, in which the decision is made by a dictator drawn from the most reliable group; expert unanimity, whereby a group accepts a proposition iff it is accepted by all members of the most reliable group; and expert majority rule, whereby a proposition is accepted iff it is accepted by a majority of members of the most reliable sub-group. As the epistemic performance of these types of rules has already been compared and unanimity rules shown to produce poor group positive reliability, we will confine attention to expert dictatorship and expert majority rule.

There are as many variants of the second category as there are ways of placing weights and fixing thresholds. Given the independence assumptions made here and on the assumption that the utility of true and false positives and negatives are symmetrical, it is known that the optimal weights to use are the competencies of the individuals, or for mathematical convenience, the logs of these competencies, with the threshold being simply the point at which the weighted sum of votes for a proposition exceeds those against (see Grofman et al (1983), Ben-Yari and Nizan (1997) and Dietrich (2006)).

In the very simple implementation that we explore here, we assume that individuals’ positive reliabilities r_j^i are identical to their negative reliabilities \bar{r}_j^i , so that $c_j^i = r_j^i / (1 - r_j^i)$, and that there is a (roughly) normal distribution of competence in the population which we then approximate by dividing the population into five ‘reliability-pentiles’ of individuals of the same reliability, with the reliability of each pentile being determined by the parameters of the background normal distribution. The important observation to make is that this modelling is intentionally crude. Any distribution of reliabilities in a group is sufficient for our result provided that it is heterogeneous (to allow for the exploitation of epistemic differences between individuals) and symmetric about the mean (so that the Condorcet Jury Theorem results hold³). The aim is to bring

³Grofman, Owen and Feld (1983) Theorem V states that if the distribution of competencies is symmetric then we obtain results analogous to the classic Condorcet Jury Theorem by substituting average reliability for homogeneous reliability. What is precisely meant by a symmetric distribution and by analogous results is unclear. We were unable to find a published proof of the full Condorcet Jury Theorem for heterogeneous reliabilities. Owen et.al. (1989) shows that the asymptotic Condorcet Jury Theorem, which states that in the limit the probability of a correct majority verdict tends towards certainty, holds for heterogeneous reliabilities irrespective of their distribution. It therefore seems plausible that the non-asymptotic Condorcet Jury Theorem, which states that the probability of a correct majority verdict is monotonically increasing with group size, also applies to cases with heterogeneous reliabilities. That the results of a homogeneous and heterogeneous group are analogous (and not identical) seems important. Some groups of voters may have identical average reliabilities but different group reliabilities depending on the distribution of reliabilities, as the following example shows:

$$\begin{aligned} (r_J^i, r_J^j, r_J^k) &= (0.6, 0.6, 0.6), r' = 0.6, r_J^G = 0.648 \\ (r_J^i, r_J^j, r_J^k) &= (0.3, 0.6, 0.9), r' = 0.6, r_J^G = 0.666 \end{aligned}$$

out important characteristics of voting rules in populations with heterogeneous competence, rather than to realistically describe actual distributions of competence (which anyway will vary from issue to issue, population to population). We assume that the roughly normal distribution of reliability is identical across all propositions, so an individual with low reliability on the first proposition will have identically low reliability on the remaining propositions. Similarly a different individual with high reliability on the first proposition will have identically high reliability on the remaining propositions.

Let the reliability of individuals in the i th reliability-pentile be r^i , with r^* being the upper bound of the reliabilities, and \bar{r} representing the average reliability of agents in the group. Then these assumptions imply:

1. *Expert Dictatorship*: $r_j^G = r^*$.
2. *Majority Rule*: As before the group accepts a proposition if and only if a majority of individuals in the group vote for it. Now, however, individual reliability varies, so group reliability on any proposition is given by⁴:

$$r_j^G = \sum_{S \subset N: |S| > \frac{n}{2}} \prod_{i \in S} r_j^i \prod_{i \notin S} (1 - r_j^i) \quad (3)$$

3. *Expert majority rule*: The group accepts a proposition only if the majority of the k most competent individuals of the group vote for it. Given the assumption of homogeneous reliability within pentiles this means that:

$$r_j^G = \sum_{h > \frac{k}{2}} \binom{k}{h} r^{*h} (1 - r^*)^{k-h} \quad (4)$$

4. *Competence weighted majority rule*: The group accepts a proposition only if the weighted sum of votes for it exceeds the weighted sum against, with the weight on the vote of any individual voter i being given by $\ln(c^i)$. This means that:

$$r_j^G = \sum_{S \subset N} \prod_{i \in S} r_j^i \prod_{i \notin S} (1 - r_j^i)$$

where the sum is taken over all sets $S \subset N$ such that $\sum_{i \in S} \ln(c^i) > \sum_{i \notin S} \ln(c^i)$ ⁵

We have assumed that the reliability of individuals does not vary by proposition and that the probability that they vote on any premise proposition is independent of whether they vote on they others. Hence application of a premise-based procedure in combination with any of four rules gives:

$$r_0^G = (r_j^G)^5$$

The monotonicity of the non-asymptotic Condorcet Jury Theorem holds for groups with heterogeneous reliabilities provided two conditions hold: firstly, that we retain current group members as group size increases; and secondly that the average reliability remains constant as group size increases. The first of these conditions is an institutional design feature we can choose to employ. By the law of large numbers, the second condition becomes increasingly likely as group size increases.

⁴See Owen et.al. (1989). According to Grofman et.al. (1983) Theorem V the probability of a correct majority winner given a symmetric distribution of competencies can be approximated with the following

formula: $r_j^G = \sum_{h > \frac{n}{2}} \binom{n}{h} \bar{r}^h (1 - \bar{r})^{n-h}$

⁵Adapted from Owen (1989).

Once again we can display the group competencies induced by these rules on single and multiple propositions for a range of values for mean positive reliability r and population size n .

Table 3: Heterogeneous Reliability, $r' = 0.49$

n		Expert Dictator	Majority Rule	Expert MR.	Weighted MR.
5	<i>SP</i>	0.86	0.48	0.86	0.92
	<i>MP</i>	0.46	2.51×10^{-2}	0.46	0.65
15	<i>SP</i>	0.86	0.47	0.94	0.99
	<i>MP</i>	0.46	2.12×10^{-2}	0.75	0.96
105	<i>SP</i>	0.86	0.40	1	1
	<i>MP</i>	0.46	1.06×10^{-2}	1	1
1005	<i>SP</i>	0.86	0.23	1	1
	<i>MP</i>	0.46	5.82×10^{-4}	1	1
10005	<i>SP</i>	0.86	0.01	1	1
	<i>MP</i>	0.46	6.41×10^{-11}	1	1

Table 4: Heterogeneous Reliability, $r' = 0.51$

n		Expert Dictator	Majority Rule	Expert MR.	Weighted MR.
5	<i>SP</i>	0.88	0.52	0.88	0.92
	<i>MP</i>	0.51	0.04	0.51	0.65
15	<i>SP</i>	0.88	0.54	0.96	0.99
	<i>MP</i>	0.51	0.04	0.80	0.96
105	<i>SP</i>	0.88	0.60	1	1
	<i>MP</i>	0.51	0.08	1	1
1005	<i>SP</i>	0.88	0.77	1	1
	<i>MP</i>	0.51	0.28	1	1
10005	<i>SP</i>	0.88	0.99	1	1
	<i>MP</i>	0.51	0.96	1	1

These studies show that there can be significant epistemic advantages, for a range of populations sizes, to be gained from concentrating decision making powers in the hands of a subset of the population. As noted earlier, Condorcet's Jury Theorem shows that if reliability is greater than a half then as the number of voters tends towards infinity, simple majority rule will generate a group reliability of certainty. But what happens at the limit is less interesting than the reliabilities for realistic group sizes. The results above show that where the average reliability of voters is greater than a half and group size ranges between 5 and 10,000 then both expert majority rule and weighted majority rule are vastly epistemically superior to simple majority rule. In fact group size need only be around 100 for these aggregation functions to generate group reliability close to certainty.

Equally significant are the results for when average individual reliability is less than a half. In this case, under majority rule group reliability worsens as group size increases. In contrast under both expert majority rule and weighted majority rule group reliability increases with group size. In fact, once again a group size of approximately 100 voters is sufficient for these aggregation functions to generate group reliability close to certainty. The mechanism behind this result is simple enough. The ability

to ignore the least competent group members (in the case of expert majority rule) or discount the least competent (in the case of weighted majority rule) is epistemically virtuous, albeit at the cost of equality of participation. The explanation for the slight epistemic superiority of weighted majority rule over expert majority rule is that in the latter the judgements of the least competent voters are not simply ignored, but treated as negative indicators. For example if a particularly incompetent voter votes ‘false’ on a proposition, the mechanism of weighted majority rule might assign one or more votes in favour of it.

Although rules exploiting competence heterogeneity have epistemic advantages, the implementation of them can be highly problematic. How is it to be determined, for instance, who are the most competent? Societies do of course have mechanisms for settling this question in different domains: educational qualifications and experience being perhaps the two important routes to becoming a recognised expert in a field. Decentralisation of judgements to experts in these ways is not generally perceived to be contrary to democratic principles of equal participation, primarily because these judgements are inputs to the political system, not outputs. When public health policy is under political review, for instance, expert advice will normally be sought from health care experts, but this advice is not automatically translated into policy. Rather it is weighed up by voters or, more typically, their representatives, and synthesised with other kinds of judgements (regarding policy goals, resources, etc..) before a policy is decided upon. When expert judgement plays this role it is more appropriately regarded as part of the deliberative process that precedes voting, than as a feature of the aggregation process itself.

It is in the policy arena that the problem of implementation is most acute. Rule by an expert oligarchy would clearly be at the expense of the ideal of equal participation in the political process. One possible way around this problem is to settle the question of who are the most competent by aggregation of individual opinion. In this domain, as in all others, some rules will fare better than others in some conditions, and worse under others. In particular, if individuals are minimally competent at judging the competence of others on a particular proposition, then majority rule will fare well as a means for picking the most competent to decide on the truth of *P*. In general if individuals are better at judging competence on *P* than at judging *P* itself, then (practical considerations aside) there will be benefit attached to using majority rule to settle the question of who should play the role of expert in the implementation of an expert majority rule verdict on *P*. Furthermore, doing so will allow for equal, indirect participation in decision making, since all will have any equal say in whose opinion is to carry the most weight in deciding on *P*.

Whether it is plausible that individuals have such second-order reliability with regard to other agents will depend on the issue at hand. At the end of the paper, we will give an argument for why it should not be assumed that individuals do generally have such knowledge. For the moment the important point to note is that once we allow for heterogeneity in individual reliabilities with regard to these second-order judgements about first-order reliability, then the argument in favour of weighted majority, and against simple majority rule, will apply again. And this in turn will raise then problem of equality of participation in decisions about whose judgements should carry much weight. Turning to a third-order aggregation of individual judgements about second-order individual judgements about first-order ones is not only impracticable, but threatens an infinite regress. In other words, once there is heterogeneity ‘all the way up’, we seem to be stuck with a conflict between the epistemic virtue of high group reliability and the procedural virtue of equal participation.

The elephant in the room in this discussion is however the fact that reliability is not uniform across propositions, so any procedure for settling who are the relevant experts will need to be of a general purpose kind that can be routinely employed whenever decisions have to be made. The enormous practical difficulties associated with organising a vote to determine the assignment of reliabilities to individuals every time a decision needs to be made on an issue, more or less rules this out as such a general purpose scheme. On the other hand if reliability was transparent, having reliability weights set by a committee would be a practical way of dealing with the problem. The problem is that in conditions of anything less than complete transparency, this system is very open to abuse and is likely to lack legitimacy. It would be better if such knowledge about relative reliability could be used by the individuals themselves. It is this possibility that we now consider.

4 Proposition-relative reliability

In our final treatment we consider a model in which individuals have both varying degrees of reliability on different propositions and that they know what these reliabilities are. We will not assume that individuals know each other's reliabilities. When reliability varies from proposition to proposition a second dimension of gains from specialisation opens up, namely those that derive from individuals specialising on propositions that they are most reliable to judge. Once again there are various possible implementations of this idea including allowing individuals to vote only on those propositions for which their reliability is above some threshold and then have the group accept propositions iff either all or a majority of those that qualify to vote, do in fact vote for it. But these rules too are unsatisfactory in virtue of the fact that they disenfranchise individuals to some degree.

Instead we suggest that allowing individuals to decide whether or not to vote, or where to concentrate their votes, offers a way of realising the gains from specialisation without contravening the principle of equal participation. Individuals themselves must be trusted to decide whether or not, or to what degree, their casting a vote on some issue is truth-conducive in the sense of improving group reliability. The simplest way of doing this would be to allow individuals to refrain from voting on any issue for which their reliability is below some threshold - 0.5 for instance. A more sophisticated treatment would allow individuals to attach more votes to those propositions on which they are most reliable. The rule we study here, which we dubbed multiple-vote majority rule, does just this.

Under multiple-vote majority rule, every voter is allocated the same number of votes, proportional to the number of propositions on which collective judgement must be reached. Individuals can then place as many of these votes as they wish on each proposition provided that the total votes cast does not exceed the number allocated. A proposition is collectively accepted if and only if the number of votes for it exceeds the number against. In our simple model for the implementation of this rule, we assume that the distribution of reliability for each proposition is exactly as it was in the previous case, with the average reliability of voters being the same across all five propositions. However the identity of the voters at each reliability level varies between propositions. We can then compare the following rules under this treatment in terms of the group reliability that they yield.

1. *Expert majority rule*: Given our simplifying assumptions, group reliability on

both a single proposition and on the set of them is the same as the previous case.

2. *Competence weighted majority rule*: The group accepts a proposition only if the weighted sum of votes for it exceeds the weighted sum against, with the weight on the vote of any individual voter i being given by $\ln(c_j^i)$, as in the previous case. Hence group reliability on both a single proposition and on the set of them is the same as the previous case.
3. *Multiple-vote majority rule*: The group accepts a proposition iff the majority of votes are in its favour. Given our simplifying assumptions this means that group reliability on any proposition will depend on individual reliability in exactly the same way as it does for majority rule. The reliabilities themselves will be different, however, since the majority will contain more votes from competent individuals. The consequences of employing multiple-vote majority rule depend on how individuals choose to distribute their votes over propositions. We consider two rules:

- i *Relative reliability*: Individuals assign to the j th proposition the percentage of their votes equal to:

$$\frac{r_j^i}{\sum_{j=1}^5 r_j^i}$$

- ii. *Relative competence*: Let $J^{i*} \subseteq J$ be the set of all propositions with respect to which individual i has a competence greater than one (i.e. on which i is at least minimally competent). Then for all $P_j \in J^{i*}$, $c_{j*}^i > 1$ individuals assign to the $P_j \in J^{i*}$ the proportion of their votes equal to:

$$\frac{\ln(c_j^i)}{\sum_{j \in J^{i*}} \ln(c_j^i)}$$

How then do these rules compare in terms of the group reliability they induce? In the table below, is displayed the probabilities for each of the rules of delivering the correct verdict on a single propositions and on the whole set of them, for a range of values for mean positive reliability r and population size n .

Table 5: Propositional-relative Reliability, $r' = 0.49$

n		Expert MR	Weighted MR	MVMR -i.	MVMR -ii
5	<i>SP</i>	0.86	0.92	0.78	0.86
	<i>MP</i>	0.46	0.65	0.28	0.46
15	<i>SP</i>	0.94	0.99	0.86	0.95
	<i>MP</i>	0.75	0.96	0.48	0.79
105	<i>SP</i>	1	1	1	1
	<i>MP</i>	1	1	1	1
1005	<i>SP</i>	1	1	1	1
	<i>MP</i>	1	1	1	1
10005	<i>SP</i>	1	1	1	1
	<i>MP</i>	1	1	1	1

Table 6: Propositional-relative Reliability, $r' = 0.51$

n		Expert MR	Weighted MR	MVMR -i.	MVMR -ii
5	<i>SP</i>	0.88	0.92	0.81	0.87
	<i>MP</i>	0.51	0.65	0.34	0.51
15	<i>SP</i>	0.96	0.99	0.90	0.97
	<i>MP</i>	0.80	0.96	0.58	0.84
105	<i>SP</i>	1	1	1	1
	<i>MP</i>	1	1	1	1
1005	<i>SP</i>	1	1	1	1
	<i>MP</i>	1	1	1	1
10005	<i>SP</i>	1	1	1	1
	<i>MP</i>	1	1	1	1

All three voting systems perform well for the chosen parameters, with weighted majority rule doing slightly better than the other three. In a sense weighted majority rule represents the upper limit in terms of group reliability that can be achieved given that reliability is tied to agents, and we cannot shift the signals among agents. As long as there is at least one agent such that their reliability is not equal to 0.5 (that their reliability is either strictly greater than or strictly less than 0.5), then weighted majority rule will always deliver a social choice that is likely to be correct.

What is most striking about these results is the fact that the multiple-vote majority rule, in both versions, does very well in terms of epistemic performance. In fact for group sizes from 105 up, multiple-vote majority rule performs as well epistemically as both expert majority rule and weighted majority rule and slightly better than then former. But at the same time, multiple-vote majority rule preserves the important procedural rule of equal participation (while expert and weighted majority rule don't). This gives support to the claim that multiple-vote majority rule achieves an optimal balance between epistemic performance and respect for equality.

It is worth considering the difference between the relative reliability versus relative competence treatment of multiple-vote majority rule. The relative competence treatment has a slight epistemic advantage over the relative reliability treatment. The relative reliability treatment involves agents concentrating their votes on those propositions that they are most competent on. They will still be able to cast all of their votes, even if their competence is below a half. By contrast, the relative competence treatment involves agents ignoring all those propositions in which their reliability is below a half (i.e. their competence is less than one) and then concentrating their votes on the propositions on which they are most competent. As such, if it were the case that an agent did not have a competence greater than one on *any* proposition then they would decide (for the good of the group competence) to abstain from casting any votes. Relative competence multiple-vote majority rule does not preserve the procedural rule of equal participation. This problem did not arise in the sample calculations above as it was assumed that each agent had at least one proposition on which they were minimally competent.

The epistemic success of all these rules derives from two factors. Firstly from removing low reliability voters from a vote on a given proposition and, secondly, from giving extra weight to the votes of high reliability voters.⁶ The net effect of the two factors is to raise the average reliability of each vote cast for a proposition. But the way

⁶See Theorems II and XIII of Grofman, Owen and Feld (1983).

these factors are put to work is different in the various rules. In the case of multiple-vote majority rule they work by individuals *choosing* to abstain from an issue or to place additional votes on it. In the case of weighted majority it is by individuals' votes being (involuntarily) discounted or inflated by the procedure itself. This virtue of multiple-vote majority rule, namely that it lets individuals make the choices, is also a potential weakness. For when individuals misjudge their relative reliabilities on issues, the rule will result in false propositions being accepted with a higher probability than, say, majority rule. But it does not seem plausible on evolutionary grounds that this kind of misjudgment could be the norm. Nor is there good reason to believe that there are reliable alternative ways of weighting individuals votes.

5 Concluding Remarks

The claim of this paper has been that, under the assumed conditions of truth-motivated independent voters with some knowledge of their own relative competence on different propositions, multiple-vote majority rule achieves a better balance of epistemic performance and equality of participation under conditions of heterogeneous competence than several well-studied rules, including majority rule and weighted majority rule. To conclude, we consider an interesting objection to this claim. The objection is that the three assumptions - that agents' know their own relative competencies, that they are motivated only by the desire to reach the truth and that they do not know the relative competencies of others - are in tension with one another. For if individuals were really so motivated, they would inform each of their relative competencies and agree to a voting rule which exploited this (now common) knowledge of differences in competence amongst individuals. In particular they would agree to use of weighted majority rule, because this rule is epistemically optimal in conditions of heterogeneous competence.

There are two things to be said about this objection. Firstly, when there is common knowledge of both the judgements each individual makes on the various propositions under consideration and their competence on these propositions, then something close to a consensus can be expected to form without need for voting of any kind. For individuals will adjust their own judgements on these propositions in the light of the information they receive about other's judgements.⁷ For instance, if you truthfully report that you believe that it will rain tomorrow and that your competence on this question is very high, then I, truth-seeker that I am and in the knowledge of my own lower competence on matters meteorological, will defer to your judgement to some (and probably large) degree. But if this is right, then the issue of what voting rule to use doesn't arise (because any unanimity-preserving one will do). We only need to vote on questions when there are limits to what we know, including about the extent of what others' know.

Secondly, the objection assumes that individuals have a very rich information about their own competencies, much richer than we think is normally justified. It is one thing to assume that individuals can make essentially ordinal comparisons of their abilities on different domains and quite another to suppose that they can make essentially cardinal interpersonal comparisons of ability. In order to implement multiple-vote majority rule, it suffices that individuals can make ordinal comparisons of ability and use these comparisons to place more votes where they are more competent. On the other hand to implement weighted majority rule we need cardinal and interpersonally comparable

⁷See Bradley (2006) for a detailed examination of this claim.

information about individuals' competencies, because we need to be able to say to what degree one person is more reliable than another. And although it is true that under the assumed conditions individuals should communicate what they know, if all they know is how competent they are on one issue relative to another, all this sharing of information will not result in a situation of common knowledge of how competent individuals are relative to each other. Hence it cannot be expected that a consensus on what weights to assign votes will form.⁸

6 Computational Appendix

Here we set out some of the basic details of how the sample calculations in this paper were undertaken. Group competence for majority rule, expert majority rule, weighted majority rule and multiple-vote majority rule were calculated using a model in the Mathematica program. Given the limits of computing power and time, the model employed a Monte Carlo simulation in which votes were pseudo-randomly drawn from a pool, with the expected value of the draw identical to the value of the reliability.

For groups with heterogeneous reliabilities we assumed that the distribution of reliabilities in our group of voters corresponded roughly to a truncated normal distribution. To contrast cases in which the reliability assumption of the Condorcet Jury Theorem was / was not fulfilled we chose two different means: $\mu = 0.49$ and $\mu = 0.51$. For ease of calculation we split this distribution of reliabilities into five discrete pentiles, as follows:

	Group 1 ($\mu = 0.49$)	Group 2 ($\mu = 0.51$)
1st pentile	0.1248	0.1448
2nd pentile	0.2791	0.2991
3rd pentile	0.4900	0.5100
4th pentile	0.7009	0.7209
5th pentile	0.8552	0.8752

On a given proposition a fifth of group 1 would then have a reliability of 0.1248, a further fifth of group 1 would have a reliability of 0.2791, and so on. For a given individual in group 1 there will be one proposition on which they would have a reliability of 0.1248, a further proposition on which they would have a reliability of 0.2791, and so on.

Under the voting rule of multiple-vote majority rule (MVMR), voters decide for themselves on what propositions to cast their votes. We investigated two rules for how individuals choose to distribute their votes over propositions, namely relative reliability and relative competence. Given the specific assumptions of reliabilities, these rules generate the following attribution of five votes:

Group 1 ($\mu = 0.49$)

	0.1248	0.2791	0.4900	0.7009	0.8552
Votes under majority rule	1	1	1	1	1
Relative reliability MVMR	0	1	1	1	2
Relative competence MVMR	0	0	0	2	3

⁸We are grateful to Franz Dietrich, Kai Spiekermann and Olivier Roy for their helpful and challenging comments on this paper.

Group 2 ($\mu = 0.51$)

	0.1448	0.2991	0.5100	0.7209	0.8752
Votes under majority rule	1	1	1	1	1
Relative reliability MVMR	0	1	1	1	2
Relative competence MVMR	0	0	0	2	3

Although these allocations of votes come from very specific calculations, the very same allocations could be obtained by voters using much simpler heuristics to decide where to place their votes. For example, if voters were aware that there was one proposition on which they were least reliable and one proposition on which they were most reliable, they could improve their epistemic contribution to the group by abstaining on the former proposition and casting an extra vote for the latter.

References

- [1] Anderson, E., (2006), "The Epistemology of Democracy.", *Episteme: Journal of Social Epistemology* 3:8-22.
- [2] Austen-Smith, D., & Banks, J., (1996) "Information Aggregation, Rationality and the Condorcet Jury Theorem". *American Political Science Review*, 90: 34-45.
- [3] Ben-Yashar, R. & Nitzan, S. (1997) "The optimal decision rule for fixed-size committees in dichotomous choice situations: The general result". *Int Econ Rev* 38: 175-186.
- [4] Bradley, R. W. (2006) "Taking Advantage of Difference of Opinion", *Episteme*, 3 (3): 141-155
- [5] Cohen, J. (1986) "An Epistemic Conception of Democracy", *Ethics*, 97:26-38
- [6] Dietrich, F. (2006) "General representation of epistemically optimal procedures" *Social Choice and Welfare* 26: 263-283
- [7] Estlund, D. (1997) "Beyond Fairness and Deliberation: The Epistemic Dimension of Democratic Authority" pp.173-204 in Bohman, J. and Rehg, W. *Deliberative Democracy*, Cambridge, Mass. MIT Press
- [8] Grofman, B. , Owen, F. & Feld, S.L. (1983) "Thirteen Theorems in Search of the Truth", *Theory and Decision*, 15: 261-278
- [9] Kruger, J. and Dunning, D. (2009) "Unskilled and Unaware of It: How Difficulties in Recognising One's Own Incompetence Lead to Inflated Self-Assessments", *Psychology*, 1: 30-46.
- [10] Owen G., Grofman, B and Feld, S.L. (1989) "Proving a Distribution-Free Generalization of the Condorcet Jury Theorem", *Mathematical Social Sciences*, 17: 1-16
- [11] List, C. and Goodin, R.E. (2001) "Epistemic Democracy: Generalizing the Condorcet Jury Theorem", *Journal of Political Philosophy*, 9: 277-306

- [12] List, C. (2008) Distributed Cognition: A Perspective from Social Choice Theory, in M. Albert, D. Schmidtchen and S. Voigt (eds.), *Scientific Competition: Theory and Policy*, Conferences on New Political Economy vol. 24, Tuebingen (Mohr Siebeck), 2008