

REPRESENTATION IN MODELS OF EPISTEMIC DEMOCRACY

PATRICK GRIM, AARON BRAMSON, DANIEL J. SINGER,
WILLIAM J. BERGER, JIIN JUNG AND SCOTT E. PAGE
zberger@sas.upenn.edu

ABSTRACT

Epistemic justifications for democracy have been offered in terms of two different forms of information aggregation and decision-making. The Condorcet Jury Theorem is appealed to as a justification in terms of votes, and the Hong–Page ‘diversity trumps ability’ result is appealed to as a justification in terms of deliberation in the form of collaborative search. Both results, however, are models of full and direct participation across a population. In this paper, we contrast how these results hold up within the familiar structure of a representative hierarchy. We first consider extant analytic work that shows that representation inevitably weakens the voting results of the Condorcet Jury Theorem. We then go on to show that collaborative search, as modeled by Hong and Page, holds its own within hierarchical representation. In a variation on the dynamics of group search, representation even shows a slight edge over direct participation. This contrast illustrates how models of information aggregation vary when put into a representative structure. While some of the epistemic merits of democracy are lost when voting is done hierarchically, modeling results show that representation can preserve and even slightly amplify the epistemic virtues of collaborative search.

INTRODUCTION

Beyond their ethical and social virtues, epistemic democracy theorists argue that there are distinctly epistemic advantages of democratic systems. When there is a correct answer to a question, democracies are more likely to get it right than the alternatives because of their superior ability to aggregate and process information. Two of the ways democratic systems aggregate and process information are by voting among alternatives and formulating and sharing policies and information through deliberation. Hélène Landemore characterizes these two forms of democratic information pooling as ‘counting’ and ‘talking’ (Landemore 2013). Elizabeth Anderson speaks of ‘votes’ and ‘talk’ (Anderson 2006).

Two formal results about these ways democracies deal with information have been used in epistemic arguments for democracy: the Condorcet Jury Theorem (Marquis de Condorcet 1785 [1995]) and the Hong–Page diversity result (Hong and Page 2004). A number of arguments appeal to the Condorcet result as a vote-based epistemic justification for democracy (Cohen 1986; Miller 1986; Grofman and Feld 1988; Ladha 1992; Estlund 1994, 2008; Gaus 1997; List and Goodin 2001). More recent work has appealed to the Hong–Page results as a model of group deliberation or group problem solving

(Page 2007, 2011; Landemore and Elster 2012; Landemore and Page 2015). Several recent sources cite both Condorcet and Hong–Page as epistemic justifications for democracy (Anderson 2006; Landemore 2013; Schwartzberg 2015).

These epistemic justifications of democracy in terms of either votes or deliberation tend to assume full and direct participation across the population.¹ But, of course, such a ‘pure’ democratic form is a distant abstraction from almost all contemporary instantiations. Present-day political democracy is almost invariably *representative* democracy, in which democratic participation is filtered through political institutions with varying levels of hierarchical representation (Dahl 1956; Lijphart 1999; Bednar 2009). While we do have reservations about using Hong–Page as a full model of group discussion, deliberation, or ‘talk,’ other substantial contributions to this literature have taken it as such. We see Hong–Page first and foremost as a model of collaborative search across an epistemic landscape. This process too often exhibits a hierarchical structure when used by democratic institutions as a means of gathering information and crafting policy (Rittel 1965; Allen and Cohen 1969; Anderson and Brown 2010; Cantimur 2016). So while not all deliberation is this form of collaborative search, to align with the extant literature, we’ll use the terms “discussion” and “deliberation” to pick out this kind of collaborative search.

In this paper, we investigate whether the benefits claimed for open and democratic participation carry over to hierarchical representation by comparing the effect of representation in the two formal models that have been appealed to. In doing so, we find a contrast in the information aggregation processes at work in the Condorcet Jury Theorem and the Hong–Page diversity model. When it comes to votes, as modeled by the Condorcet theorem, representative democracy does worse in terms of accuracy than its purer simple-majority relative. But, adding representation to the Hong–Page model does not compromise the epistemic virtues of collaborative search in that model. As we show, for a range of variations, the virtues of diverse collaborative search hold up well within a representational structure. For one particular instantiation of deliberative dynamics, representational hierarchy even does slightly better than direct participation. While some of the epistemic merits of democracy are lost when voting is done hierarchically, our results indicate that representation can preserve and even, in some cases, robustly amplify the epistemic virtues of collaborative search.²

In the first section, we outline the downside of representative voting in the Condorcet result, and in an appendix, we give a new and more direct proof of an important relevant result from Boland (1989). In later sections, we contrast this with the results of computer simulations that show the resiliency of representative *search* as conceived of in the Hong–Page model. Throughout, we offer formal results from agent-based models using abstract structures to model representation and abstract dynamics to model voting and collaborative search. As is well known, the structure of ‘voting’ in the Condorcet theorem abstracts from real voting in many regards, with voter independence as a major assumption.

1 For a notable exception see Landemore (2013).

2 On a political interpretation, our results might be taken to support appeals for various forms of non-electoral deliberation – public spheres, citizen forums, and ‘mini-publics’ (Fishkin 1991, 1995, 2009; Habermas 1996; Fung 2006; Urbinati and Warren 2008; Warren and Pearse 2008; Niemeyer 2011; Grönlund *et al.* 2013) – by showing the robustness of the epistemic virtues of democratic collaboration within representational structures.

Following in that tradition, our representational structures follow those in Boland (1989) and Boland *et al.* (1989), and they satisfy the four criteria set for political representation in Urbinati and Warren (2008). Formats for collaborative group search are borrowed directly from Hong and Page (2004), with variations considered below. Here, as always, application of formal results requires a wary eye as to whether the abstractions employed adequately capture relevant aspects of political and epistemic reality.

I. THE DOWNSIDE OF REPRESENTATIVE HIERARCHY IN THE CONDORCET JURY THEOREM

The Condorcet Jury Theorem outlines simple conditions under which a collective will be more likely to make a correct decision than an individual. For any odd number of independent voters n (obviating the problem of ties), each with probability p of getting the question right, the probability of a correct vote of a majority is:

$$\sum_{k=(n+1)/2}^n \binom{n}{k} p^k (1-p)^{n-k}$$

What the Condorcet justification for democracy points out is that, as long as voters are more likely to be right than wrong (i.e. $p > 0.5$), the total probability of the majority being right quickly becomes stronger as the population size increases. Figure 1 shows the probability of the majority being correct across increased numbers of voters for four probabilities of individual correctness.

The Condorcet result relies on a number of questionable assumptions, including that voters have identical probabilities of being correct and that their votes are probabilistically independent. Recent literature shows that both assumptions can be relaxed while maintaining the important result (e.g. Boland 1989; Boland *et al.* 1989; Ladha 1992, 1995; Dietrich and Spiekermann 2013). Here we'll focus on another assumption of those who use the Condorcet result to justify democratic institutions, namely the assumption that the result will extend to *representative* democracies.

Boland (1989) demonstrates that the Condorcet result is significantly weakened when applied in the simplest form of hierarchical representation. We offer a substantially more direct proof of that result in the appendix. To understand that result, consider a toy example with 9 individuals in a representational system. If we assume that each individual has a probability of correctness of 0.6, the probability that the majority of them is right, in a direct vote, is 0.733. Now suppose that the 9 are divided into 3 groups of 3, and each small group first votes to determine the vote of their representatives on a second level. Those representatives simply reflect the majority decision of their constituents, voting again as a group of 3. To find out the probability that a majority of 3 representatives will get the answer right, we first consider what the probability of success of the small groups is. Since there are only 3 people in the small groups, again using probability of 0.6 for each individual, the probability of a correct majority decision in a small group is 0.648.³ So, the 3 representatives then each have a 0.648 probability, meaning that

3 Goodin and Spiekermann (2012) offer some related calculations for differently structured representative bodies.

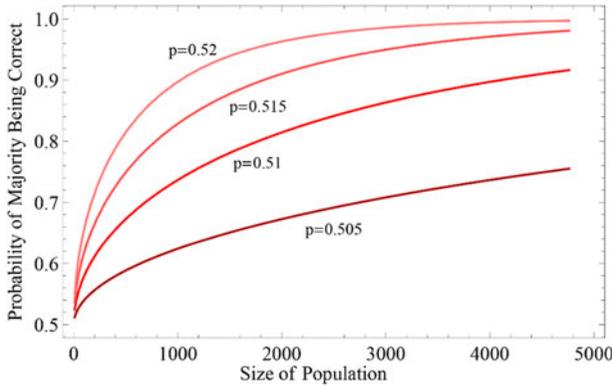


Fig. 1. Probability of a majority of different odd-numbered sizes being correct on a binary question with different homogeneous probabilities of individual members being correct.

the probability of a majority of the three representatives being right is 0.716. This is lower than the probability of the direct vote of 0.733. Hence, representative hierarchy compromises the epistemic virtues of democracy.

This result can be generalized: with n voters, a direct majority vote always has a higher probability of producing the correct outcome than a representative hierarchy in which the voters are partitioned into g (odd) subsets each of which has an odd (and not necessarily equal) number of voters.⁴ For the specific case in which there are g groups of g people each, Figure 2 shows comparative probabilities of a correct decision from a representational structure and for a simple majority decision of all $n = g \times g$ individuals involved.⁵

As the population size goes to infinity, probabilities for both representational and simple majority decisions go to 1. But for any finite n , the probability of a correct decision with a straight democratic vote will be higher than a decision within the comparative representative structure (Boland 1989 and Appendix). What this result shows is that, when conceived of in this simple way, representative structures diminish the overall probability that the institution gets the question right. This is admittedly a ‘knife-edge’ result, highly dependent on assuming that the representatives’ probabilities of a correct decision match the joint probabilities of the subgroup they represent. Small variations can dramatically alter the overall probability of correctness of the group, and a greater probability of correct decisions on the part of representatives can compensate for the loss inherent in the

4 The formula for probabilities in the general case of g groups of m individuals, with both g and m odd, is:

$$r = \sum_{k=(g+1)/2}^g \binom{g}{k} q^k (1-q)^{g-k} \text{ where } q = \sum_{k=(m+1)/2}^m \binom{m}{k} p^k (1-p)^{m-k}$$

(q is the probability of a group getting the correct answer, itself calculated for each group in terms of individual probabilities p .) For even numbers, the result holds with either a calculation altered to break ties or requiring a strict $(m/2) + 1$ correct votes.

5 Although we don’t pursue the degree of difference with different numbers of groups and members, it is worthy of note that Beisbart and Bovens (2013) find mean majority deficit results to be greatest in those cases in which $g = m$.

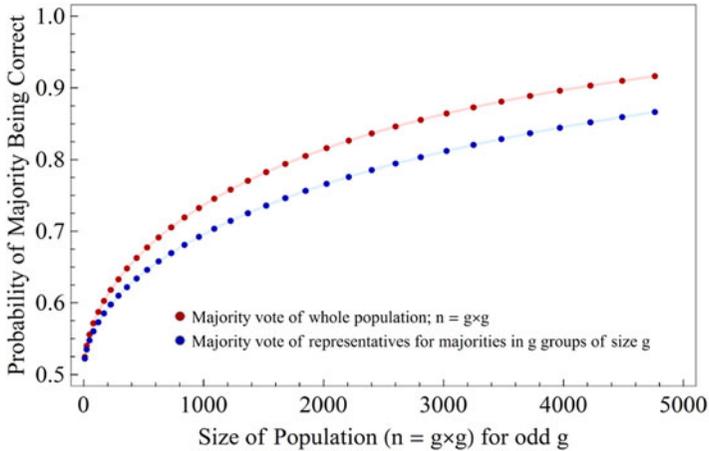


Fig. 2. Probability of being correct in a majority vote of the whole population (red) and a representative vote (blue): a majority vote among g representatives for groups of odd size g , each voting the majority vote of their constituency.

representative structure (see also List 2005). It is therefore unclear whether this formal result can be used to defend actual representative democracies.

Importantly, this doesn't ultimately undermine the Condorcet argument as a defense of representative democratic institutions. It's still true that including more voices in the political process, either as representatives or constituents, improves the outcomes of the process, epistemically speaking. Although representation does diminish some of the benefits, the epistemic costs might be justified by the increased institutional feasibility of representative voting. What's surprising is that even the demonstrated costs of representation go away when we turn to collaborative search.

2. COLLABORATIVE SEARCH AND THE HONG–PAGE RESULT

One way that groups pool information is by democratic voting. In real democracies, voting operates through representation. Another way groups pool information is by one or another form of group problem solving or collaborative search. Lu Hong and Scott Page (2004) offer a formal model for such a procedure which incorporates features of group interaction that go well beyond mere voting. It is on that basis that it has been embraced in later work as a model for democratic 'talk' as opposed to 'votes' in epistemic democracy (Anderson 2006; Landemore 2013; Schwartzberg 2015).

What Hong and Page claim to show is that in a particular form of collaborative search, diverse groups will regularly outperform groups composed of the best-performing individuals. They term this the 'Diversity Trumps Ability' result. There are actually two forms of the result. Hong and Page offer a mathematical proof for the limiting case of an infinite population, in which the diverse group performs best with probability 1 (Hong and Page 2004: 16387 ff). That proof is used to explain what happens in more realistic cases of finite populations in which the diverse group performs better with high

probability. For finite populations, the result is shown to obtain in simulations rather than in formal proof. The same is true of our work here.

Hong and Page start with an ‘epistemic terrain’ or problem space containing n positions $\{1, 2, \dots, n\}$. The positions are arranged in a circle, moving from 1, 2, ..., $n-1$ to n back to 1, 2, ... again. Each of the spots on the terrain is mapped to a real-numbered score drawn from a uniform distribution on the interval $[1, 100]$, which can be thought of as scores representing how well that ‘epistemic spot’ answers some question. Starting from any given point, agents try to find the highest points on the terrain by using their assigned heuristics defined as ordered sets of k numbers, each between 1 and l .

An example is helpful: Each agent has a heuristic, which is a set of k integers. Elements of the heuristic are used in sequence. In a case in which $n = 2000$, $k = 3$ and $l = 12$, for example, an agent’s heuristic might be (3, 8, 2). Starting at some arbitrary point on the circle, that agent will first evaluate the value of a point three spots to the right from its origin, comparing it with the value at its current position. If the value at this spot is higher than where it currently is it will move to the new position (otherwise it will stay put). It will then try its second heuristic number, evaluating the spot 8 away. If it too is higher, it will move to that point and then try its third heuristic number 2. The agent keeps checking points to the right using its heuristic numbers in turn, moving to higher points, until it reaches a stopping point at a local maximum from which none of its heuristic numbers can guide it to a point with a higher value. Figure 3 illustrates another example.

We can measure the overall expertise of an individual agent in this model in terms of how well the agent can be expected to do in finding the maximums on the terrain from random starting points. Let $\varphi(i)$ denote the stopping point for a heuristic φ from an initial starting point i , with $V[\varphi(i)]$ the achieved value at that stopping point. The overall performance of an agent with heuristic φ is the expected value $E[V; \varphi]$ of the stopping points for each of the n possible starting positions i . If we assume that all of the starting points are equally likely, that simply amounts to the average of the $V[\varphi(i)]$ for all i . Given any specific epistemic terrain, agents individuated by their heuristics can be ranked in terms of overall performance. Hong and Page’s ‘best-performing’ agents on a terrain are those with the highest expected values: those with the highest average value reached from the n points on that specific terrain.

The Hong–Page result compares the expected performance of collaborative groups of agents exploring a terrain together. The comparison is between expected values for (a) a group of ‘best-performing’ agents on a given terrain – the 9 with the highest expected values on that terrain, for example – and (b) a diverse group of 9 agents assigned random heuristics. The way that agents act together in groups is a variation on the pattern of individual exploration outlined above. The procedure is a sequential relay among, in this example, the 9 group members. Starting from a given point, the first agent finds the highest point she can get to (as outlined above). Once she has reached that personal stopping point, the baton is passed to the next member in the group. The second member of the group starts from the stopping point of the first member in order to see whether her heuristic can reach a point with a higher value from there. The relay procedure continues in sequence through the 9 participants until a point is reached from which none of the 9 agents can find a point with a higher value.

The overall group performance can be measured like individual performance. Let $g(i)$ denote the relay stopping point for a group g from an initial starting point i , with $V[g(i)]$

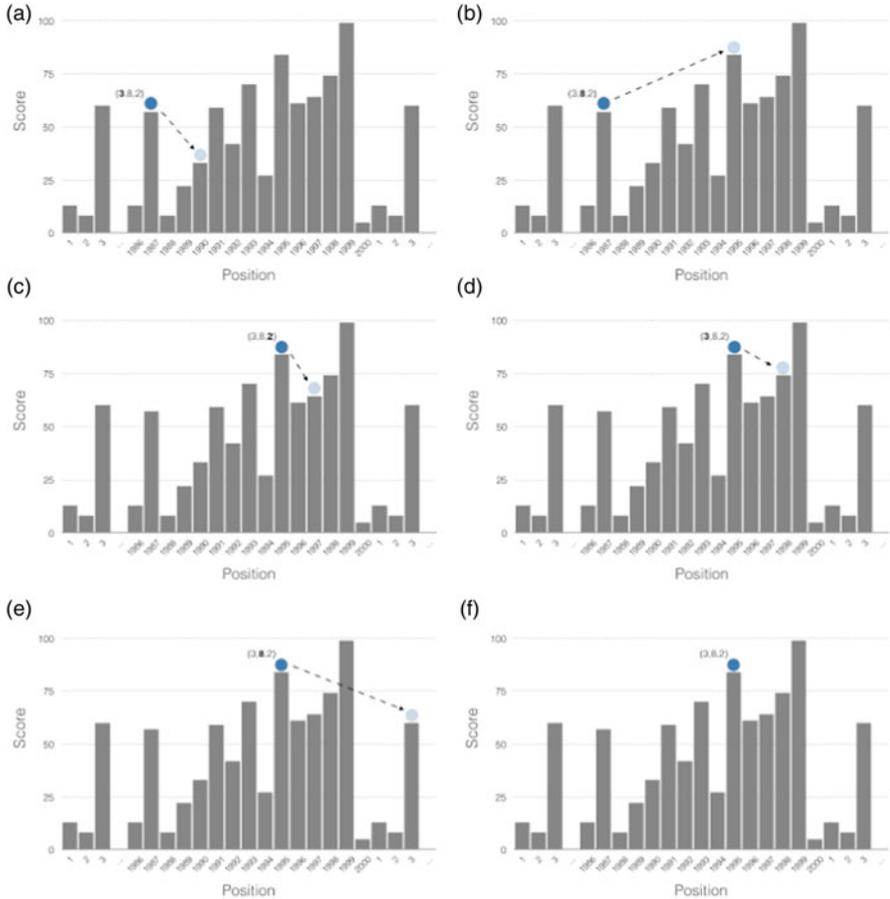


Fig. 3. In this example, the agent begins at position 1987 with the heuristic (3, 8, 2). The agent first checks the value at location 1990 (1987 + 3), but finding it lower than that of its current position it remains at position 1987. Moving to its next heuristic number 8, the agent finds location 1995 (1987 + 8) to have a higher value and so moves to the new location. From there the agent deploys its heuristic number 2, then 3, followed by 8. In checking positions 1997, 1998, and 3 (1995 + 2, 1995 + 3, and 1995 + 8 mod 2000, respectively), the agent finds no spot with a greater value than the one it currently occupies, and therefore stops.

the value achieved at that stopping point. The performance for a group g is the expected value $E[V; g]$ of the stopping points for each of the possible group starting positions i . If we assume that all of the starting points are equally likely again, that amounts to the average of the $V[g(i)]$ for all i .

What Hong and Page establish is that a group of agents with random heuristics usually does better in this relay group problem-solving format than a group of the individually best-performing agents. By the theorem, diversity trumps ability with probability 1, given stark conditions regarding group and population size, problem difficulty, and group diversity. When those conditions are relaxed, the probability is no longer 1. That said, replications show the result to be extremely robust: the epistemic success of a group of random heuristics proves consistently superior to that of a group of those

which individually score the best.⁶ This is the ‘Diversity Trumps Ability’ result, which has been taken to show the virtues of democratic discussions over those that include only the individual experts (Anderson 2006; Landmore 2013; Schwartzberg 2015).

As noted above, we have reservations interpreting the Hong–Page result as a model of group deliberation or discussion. Real deliberation is a complex social process (Stasser and Dietz-Uhler 2001; Tindale *et al.* 2001; Kohn *et al.* 2011; Levine and Smith 2013; Witte and Davis 2014; Levine and Tindale 2015), though other attempts to model deliberative group dynamics do often model it in terms of either simple preference or judgment aggregation (Austen-Smith and Feddersen 2005, 2006; Landa and Meirowitz 2009; Mathis 2011; Perote-Peña and Piggins 2015; Dietrich *et al.* 2016; see List 2018).⁷ The Hong–Page model has the virtue of capturing that aspect of democratic deliberation represented by collaborative information-sharing with an eye to the exploration of alternatives. So while it does capture the idea of information passed from one individual to another, the modeled communication between individuals does not involve reasons or arguments, is not open to challenge or cross-questioning, and has a specific one-way flow (the order of the ‘relay’).⁸ Although touted as a representation of democratic ‘talk,’ it is perhaps wiser to treat the Hong–Page model more conservatively as a model of one particular dynamic of group search. In that more conservative sense then, for a specific kind of search dynamics, we should see the Hong–Page result as indicating the virtues of diverse groups over groups of the individually best-performing agents.

3. REPRESENTATIONAL HIERARCHY: RESILIENCE OF THE HONG–PAGE RELAY MODEL

As in the case of the Condorcet theorem, we ask here how the Hong–Page model can be expanded to apply to representational hierarchies. We consider two different idealized instantiations of a deliberative hierarchy. The first uses a ‘relay’ format based on the original dynamics of the Hong–Page model. The next section uses an alternative ‘tournament’ format, with importantly different results.

In a representational hierarchy for relay discussion, deliberation is instantiated at two levels. First, the subgroups aggregate their information in the following way: Starting from an initial point on the terrain, each subgroup uses the relay format as described above to find the highest point they can achieve from that initial point. They repeat that process, treating every point on the terrain as an initial point and noting the highest place they

6 While the Diversity Trumps Ability result has been challenged on the basis of how unexpected it should be and for the relevance of the mathematical theorem offered in its explanation (Thompson 2014), we are not worried by that critique. Kuehn (2017) offers an extended response to Thompson, while Singer (Forthcoming) refutes Thompson’s critique that the result is driven by randomness rather than diversity, *per se*. It is sufficient for our purposes here, however, to note that Thompson misunderstands the work Hong and Page’s model seeks to do. Though she contends that the results of the Hong and Page model are not unexpected, the result’s mathematical obviousness is largely beside the point. The model’s import is that it tracks a good deal of what the social scientific and philosophical literature means when it discusses diversity, and provides leverage to deal with such a concept more formally.

7 The authors are currently exploring a more reason- and argument-rich agent-based model of group deliberation. In this regard see also Betz (2013).

8 In work below we consider an alternative ‘tournament’ dynamics.

can get from each of the points on the terrain. This ‘portfolio’ of their performance, which can be seen as reflecting the pooled wisdom of that subgroup, is then passed to a representative of that subgroup in the higher body. The representatives then do a similar relay deliberation, this time using the portfolios of their subgroup as their heuristic: From each initial point, the first of the three representatives reports the highest point their subgroup got to from an initial point by consulting the subgroup’s portfolio. The representative of the second subgroup advances the deliberation by starting at the point the first subgroup stopped at. The third representative then has the chance to advance the group in the same way. The rotation continues through the three representatives, stopping at a final point from which none of the representatives is able to achieve a higher value with their portfolios. The performance for the whole group is the average final value reached from each of the 2000 initial points.

As in the Condorcet parallel, a representative’s view in this model is determined by their constituents, though their view is not fixed by any particular individual view. In the Condorcet model, representatives simply pass up the majority vote of their constituents. They can thus be viewed as carrying the vote of their group up to the higher level. In the Hong–Page model, representatives are handed a portfolio from their subgroup, which gives the highest point that the group can get to from each possible starting point. In this way, representatives can be thought of as carrying up the collective knowledge of their subgroup. The representatives then combine the knowledge from their respective subgroups yielding the sum collected knowledge of the whole population. There is therefore a difference in how representation is naturally conceived of in the two models. In the Condorcet model, representatives combine the votes of their constituents, whereas in the Hong–Page model, representatives combine the search results of their constituents. Of course, this difference tracks a difference in what the two models are about: it’s votes in the first case and collaborative search (perhaps modeling aspects of ‘talk’) in the second.

We compare relay discussions among all 9 members and relay discussions among 3 representatives of 3 groups of 3 members in the Hong–Page model. In this model of discussion, in contrast to the Condorcet model of voting, representation holds its own. Results are virtually tied (see Table 1). Over a typical 1000 runs, the average scores for open and representative discussions among 9 randomly chosen individuals are 94.89775 and 94.89646 respectively – a difference of 0.0014%. Moreover, representation actually does better in terms of the absolute number of cases, which indicates just how

Table 1 Results indicate the number of runs in which the representative structure outperformed the direct set up. Top numbers in each cell indicate the number of runs in which representation beats direct, while the bottom number indicates the average margin between the two instantiations. The star indicates significance at alpha < 0.001. Note, however, that the result in the top left is not significant even at alpha = 0.05.

	Relay	Tournament
Randoms	516/1000 0.0014	1000/1000 0.1376*
‘Bests’	424/1000 -0.0293*	999/1000 0.0879*

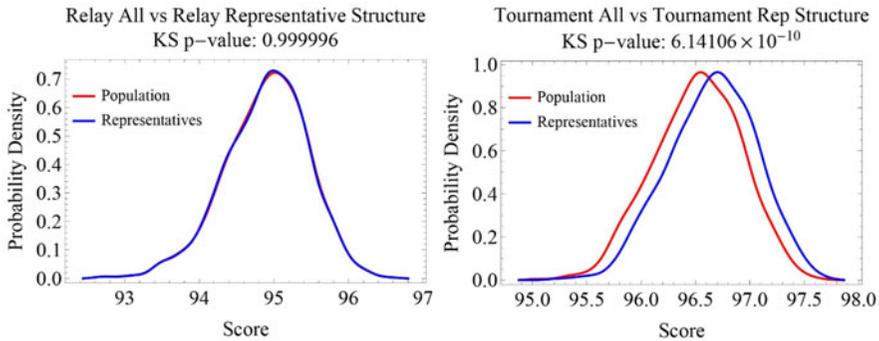


Fig. 4. The empirical probability density functions of the highest points reached by the population and representatives in the relay (left) and tournament (right) experiments. Results of the Kolmogorov–Smirnov two-sample test are shown above each.

close a tossup this relay schema is. Representation beat open discussion in 516 of 1000 cases; open discussion beat representation in 463. The left graph in Figure 4 shows nearly identical probability density functions and therefore a lack of any significant difference in terms of the Kolmogorov–Smirnov two-sample test. We get similar results when we compare an open discussion of the 9 ‘best’ heuristics with a representational hierarchy for the same individuals. Averages come in at 92.76888 and 92.74936 respectively, a difference of 0.02%.

4. REPRESENTATIONAL HIERARCHY: SUPERIORITY OF THE HONG–PAGE TOURNAMENT MODEL

Here we consider a variation on the search dynamics within the Hong–Page model, modeling group collaboration as a simultaneous tournament rather than a sequential relay. Hong and Page (2004) claim that this variation gives essentially the same results. However, with a representational hierarchy, that turns out not to be true. Within a tournament dynamic, representational hierarchy robustly does slightly better than direct participation.

In a tournament dynamic we eliminate the ‘single file’ or ‘each in turn’ assumption of the relay. For discussion among a group of 9, all 9 agents use their heuristics simultaneously, each reporting the stopping point with their highest value. The stopping point with the highest value achieved by *any* of them is then treated by all members as a new starting point. At each stage all agents in the group simultaneously leapfrog to the highest point achieved by any of them, without having to pass a baton in sequence. This format has a different dynamic and can be expected to hit different points in the process, but the final stopping point will again be a local maximum from which none of the agents’ heuristics can reach a higher point. The performance for the group is the average final value achieved from all of the 2000 starting points.

For representation we again use the simple example of 3 groups of 3. Each subgroup engages in a tournament search to create their portfolio, which lists the highest point they can get to starting at each of the 2000 points in the terrain. The representative

then uses the portfolio, which contains the group's pooled wisdom, as a guide in tournament deliberation with other representatives on the second level. On the second level, for each point on the terrain, all representatives simultaneously report the result achieved in their groups' deliberation. The point at which the highest value is reported then becomes a new starting point. Representatives then report the highest value each of their groups has reported from *that* point, and so on. The performance for the group is again the average final value achieved from all of the points.

All that has been changed in this instantiation of representation is the structure of group participation on each level: a simultaneous tournament in this case as opposed to a sequential relay in the previous section. Here, however, we have different results: this second representational structure not only preserves the virtues of group search, but amplifies them slightly. Over 1000 runs, the average value for a group of 9 random agents with tournament discussion among representatives of groups of 3 comes in higher than the average for those same 9 in open tournament discussion. The difference in each case is slight: typical averages are 96.58819 and 96.45059 respectively, a difference of 0.14%. What is most striking though is the robustness of representation's lead. However slightly, representative systems did better in *all 1000 of 1000* of the runs we did. In a tournament structure then, though not by much, representation robustly beats out direct participation.

Tournament dynamics achieve higher epistemic scores than do relay dynamics across the board, regardless of whether the discussion configuration is representative or whole-group. In the discussion above we contrasted results using 9 random agents with and without a representational decision structure. All results hold, including the better results for tournament dynamics, if we use groups of 9 'best-performing' individuals instead. Table 1 summarizes comparisons across all three variables: the two dynamics, with representation or without, and using random groups or groups of the 'best-performing.'

Our results confirm the original Hong–Page 'Diversity Trumps Ability' result and show that it is maintained in representative structures. Above, we compared the performance of groups of both random and best-performing agents separately in direct versus representative structures. The Hong–Page diversity result compares groups of random agents and groups of 'the best.' As summarized in Table 2, that result is maintained in representational structures. In more than 97% of our runs we find that a random collection of agents still outperforms a group of the best-performing individuals. The epistemic advantages of diversity remain important in the representational structure and are even enhanced by the tournament dynamic.

Figure 4 shows the outcome and significance results for relay discussion (left), in which representation is statistically indistinguishable from open discussion, and for tournament discussion (right), in which representation shows a slight edge. A Kolmogorov–Smirnov two-sample test⁹ confirms that the distribution of relay outcomes are nearly identical, but the outcome distribution for the tournaments are significantly different. Except for the fact that representation holds its own in tournament as well as relay discussion, the advantages for tournament representation are so small that one must be wary of any attempt to read off practical implications. But, the theoretically intriguing fact remains that tournament representation shows higher scores so robustly.

9 The Kolmogorov–Smirnov two-sample test was chosen because it is a non-parametric test to determine whether two samples might be pulled from the same distribution and our observed outcomes did not clearly match any parametric distribution.

Table 2 Results indicate the number of runs in which a random group of nine outperformed a group of the 9 individually best-performing agents. Top numbers in each cell indicate the number of runs in which random beats best-performing, while the bottom number indicates the average margin between the two groups. The asterisk indicates significance at $\alpha < 0.001$.

	Relay	Tournament
Direct	973/1000 2.2329*	973/1000 1.3907*
Representation	973/1000 2.1702*	977/1000 1.4404*

5. UNDERSTANDING THE REPRESENTATIONAL ADVANTAGE IN TOURNAMENT DYNAMICS

Our results indicate that groups of nine agents organized in a representational hierarchy do slightly better in a tournament dynamic than those same nine in a direct participation dynamic. Here we take some steps toward understanding that result. It cannot be that groups employing representation do better than groups with full participation on every point on a random landscape. We can construct landscapes in which either will perform better from a given point.

As a first pass, consider one way representational groups can outperform open groups. The small groups in the representational structure can sometimes get to spots that the single larger group cannot because the larger group will get stuck on a local maximum that one of the small groups, lacking a heuristic in the larger group, would not get stuck on. The small group can get past that local high point that the larger group is stuck on, managing to find a higher peak beyond. Here is an example to clarify:

Suppose we simplify individual heuristics to be single numbers (rather than triplets) of integers 1 through 9. Let the group consist of nine agents, one with each of the possible heuristic values. We'll use the terrain depicted in Figure 5: a local peak at spot 7 (value 96), a higher peak at spot 20 (value 99), and values 85, 89, and 92 at spots 5, 10, and 15. All other values are 50. Starting from point 0, the group of 9 will get stuck on 7 and won't be able to reach anything higher beyond that. But a small group with a 5, but no 7, will find the higher value at spot 20 in four hops to increasingly higher points.

That kind of set up might explain how representation-groups can do better. But there will also be cases in which the group of 9 does better than the three groups of 3. Let our groups be group A, 3 agents with heuristics {1}, {3}, and {5}; group B with {2}, {4}, and {6}; and group C with {7}, {8}, and {9}. Let our terrain be the one shown in Figure 6 with a value of 86 at spot 5, 84 at spot 6, 90 at spot 7, 82 at spots 8, 10, and 12, a high of 99 at spot 13, 94 at 15 and 50 for all other points. In that case, starting from point 0, the group of 9 will hit 90 at spot 7 with heuristic 7 and jump from there to a score of 99 at spot 13 with a heuristic 6. The smaller groups will not do so well. Group A [1 3 5] will go to 86 at spot 5, but neither 1, 3, nor 5 will take it higher. Group B [2 4 6] will go to 84 at spot 6, but neither an additional 2, nor 4, nor 6 will take it higher. Group C [7 8 9] will go to 90 at spot 7. From there it will go to 94 at spot 15 using heuristic 8, missing spot 13 in the process and with nothing but 50's from there on out.

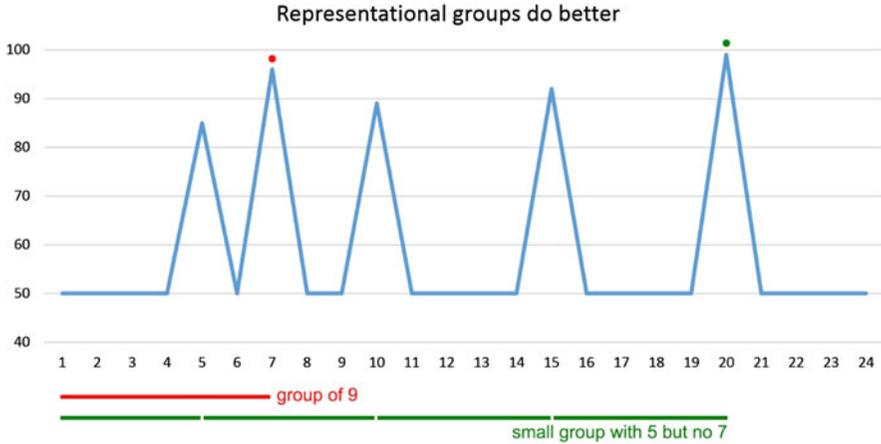


Fig. 5. A landscape on which a small group does better than the full group of 9.

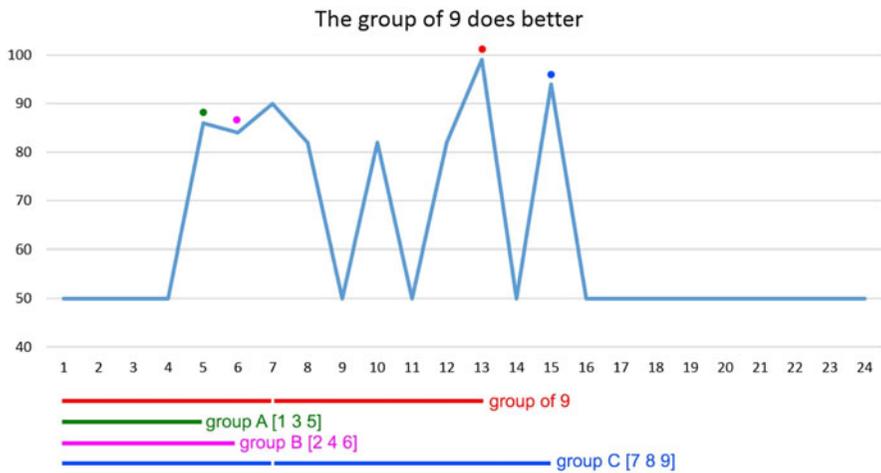


Fig. 6. A landscape on which the group of 9 does better than smaller groups.

For individual starting points on random terrains, it is therefore possible for either whole groups of 9 to do better than 3 representational groups of 3 or for the contrary to hold. In general though, they do not perform equally well. It is far more common for representational groups to do better than the full group (across 100,000 starting points, representational groups do better in 3278; the full group does better in only 183).

What explains the greater frequency with which representational groups do better? One clue is the number of points sampled by each. Over 100 runs, the average number of distinct points sampled from each of the 2000 landscape points using a tournament among 9 random agents is just over 26. The average number of distinct points sampled using a representative structure of 3 groups of 3 agents is higher at 27.5. It would make sense that a dynamic that sampled more points would occasionally have some

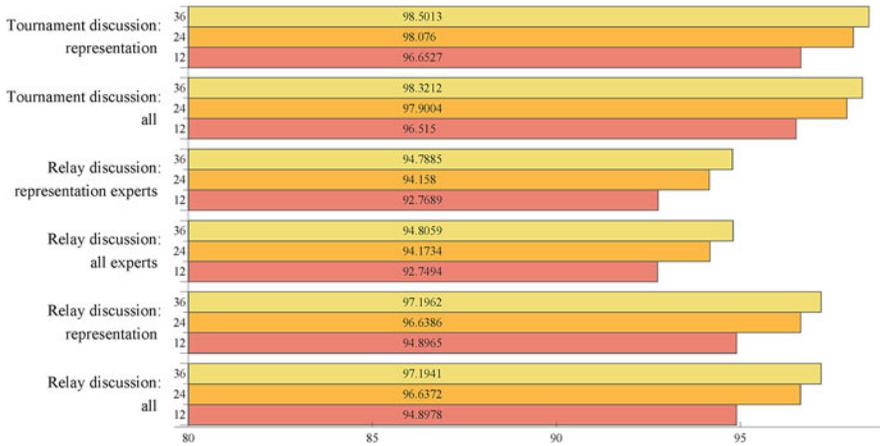


Fig. 7. The effect on performance of enlarging the pool from which heuristics values are drawn, values are scores averaged over 1000 runs.

higher points in its sample. Our hypothesis is that representational groups do slightly better because they sample a slightly larger number of points on the terrain. An explanatory hypothesis of this type is consistent with our other results as well. Tournament dynamics, which outperform relay dynamics, sample a larger number of points: an average of 27.5 points versus an average of 17.4 with representational structures, and 26 versus 17.4 with open discussion. Results from open discussion and representation are indistinguishable in the case of a relay dynamics. Both sample an average of 17.4 points.

6. ROBUSTNESS OF OUR RESULTS

With simulation results like those offered here, robustness under model variations is key to understanding how broadly the results apply. In all of the results above, we use a terrain created by picking a random number between 1 and 100 for each of 2000 spots. We might instead assign scores from a normal distribution, with, for example, a mean of 50 and a standard deviation of 10. Results will be a predictable transform of those above, but they do highlight the difference between open and representational dynamics a bit more. Using a normally distributed terrain, results for open and representational tournament discussion come in at 69.77 and 69.95 respectively rather than 96.51 and 96.65, a percentage difference of 0.25% rather than 0.14%.

Throughout, we use a model in which the three heuristic values for each agent were drawn from a pool of 12. What happens if we expand that pool, so that heuristics are composed of three numbers from a pool of 24, for example, or 36? Since the original Hong–Page effects can be understood largely in terms of redundancy (Hong and Page 2004; Thompson 2014), we should expect this to make a difference. Average results across 1000 runs for pools of 12, 24, and 36 heuristic values are shown in Figure 7.

For cases of relay representation in heuristic pools of different sizes, there is no consistency in whether the representational structure does better or worse than full participation. The same proves true in the case of best-performing individuals in full or representational

relay discussion. Whatever the size of pool heuristics, a representational structure for tournament discussion does better than full discussion among the population, though the margin of superiority declines as the pool of heuristics is increased. The central fact highlighted above, however, holds throughout: In *every* case in *every* run of 1000 the representational structure for tournament discussion comes in with a higher score.

Do results hold for larger numbers of larger groups? If we consider not 3 groups of 3 agents, but 5 groups of 5, results remain roughly comparable, with a particularly robust result for the triumph of representation in the case of tournament discussion. For relay discussion of the form considered in section III, a representational structure still does almost exactly the same as relay discussion among the population as a whole. The same holds for relay discussion among best-performing individuals. The central results that we have highlighted for tournament discussion prove robust in the expanded case: in tournament discussion representation does better (an average of 97.180 as opposed to 97.096). Here as before, the most significant fact may be that representational structure of tournament discussion once again did better in every one of the 1000 runs.

7. CONCLUSION

Epistemic democrats use formal models like those above to argue that democratic institutions provide good epistemic outcomes. Formal arguments in support of this claim have been based on models of either voting (Condorcet) or ‘discussion’ as collaborative search (Hong–Page) in a ‘pure’ or ‘direct’ democracy (Cohen 1986; Estlund 1994; Anderson 2006; Landemore and Page 2015; Schwartzberg 2015). Here, we have focused on whether these epistemic arguments for open group participation carry over to abstract models of hierarchical representation as well. Our point of departure has been to ask whether two theorems touted by epistemic democrats hold up in a representative or hierarchical structure and to compare epistemic merits in each case.

Our results indicate that the answer depends first and foremost on whether it is virtues of voting or collaborative search that are at issue. Both analysis and simulation show that the epistemic virtues of voting that appear in the Condorcet theorem are weakened in a representational structure. When it comes to votes, inclusive democracy proves epistemically superior to representational democracy, just as Rousseau might have argued (Rousseau 1762 [1978]).

The story is different for the epistemic virtues of collaborative search. Within a relay deliberation of the form that appears in the original Hong–Page model, epistemic results for representational structures prove virtually identical to results for full deliberation among all participants. In search, as opposed to voting, representation holds its own. In a simultaneous tournament structure we’ve shown that epistemic scores for a representational structure slightly, but persistently, edge out those for the population as a whole.¹⁰ Small groups search for optima independently, with those results carried up for a second-order deliberation. Under certain conditions this process proves as effective as one in which an entire population of agents seeks a best solution on the given terrain.

¹⁰ In a political interpretation, these results lend formal support to arguments for federalism (e.g. Bednar 2009) and dicta like Justice Louis Brandeis’ claim that states in a federated system act as the “laboratory of democracy.”

The high efficiency of hierarchical representation is well-established (Pitkin 1967; Dahl 1983).¹¹ With regard to epistemic benefits of collaborative search, as opposed to votes, our results indicate that representation's virtues of efficiency need not entail an epistemic sacrifice.

APPENDIX: LIMITATIONS OF THE CONDORCET JURY THEOREM IN REPRESENTATIONAL STRUCTURES¹²

Suppose there is a vote between binary exclusive options *A* and *B*, where each voter has a probability $p > 0.5$ of choosing the correct option. We compare two forms of aggregation. (1) In a direct majority, each voter casts a ballot for *A* or *B*; whichever receives the most votes is the outcome. (2) In a representative structure we assume the voters are partitioned into *g* groups of size *m*, where both *g* and *m* are odd. Voting takes place in two stages. In the first stage members of each group *g* vote using direct majority. In the second stage whichever outcome was chosen by the majority in each group is cast as the vote for that group in a second majority vote among group representatives. The option winning the second vote is the representative outcome.

With these conditions a direct majority vote is more likely than a representative structure to produce a correct outcome. This result is anticipated in Boland (1989). In this appendix we offer what we consider to be a more direct and accessible proof.

To provide insight for the general result we first consider the case of 9 voters who are partitioned into 3 groups of size 3. We can then make the following observations:

Observation 1. With 9 voters, a direct majority produces the correct outcome if and only if 5 or more of the voters choose the correct option.

Observation 2. With 9 voters partitioned into 3 groups of size 3, the representational structure always produces the correct outcome whenever 6 or more voters choose the correct option. No matter how the 6 correct votes are assigned to the 3 groups of 3, there will inevitably be a majority of groups with a majority of correct votes.

Observation 3. With 9 voters partitioned into 3 sets of size 3, the representational structure cannot produce the correct outcome if 3 or fewer voters make the correct choice. No matter how the 3 correct votes are assigned to the 3 groups of 3, it is impossible to obtain a majority of groups with a majority of correct votes.

Given these observations, the interesting cases arise when either 4 or 5 voters make the correct choice:

Observation 4. With 9 voters partitioned into 3 sets of size 3, the representative voting structure can produce either outcome when exactly 4 or exactly 5 voters make the correct choice.

To see why this is true, consider two realizations of choices and partitions using *c* for the correct choice and *i* for incorrect.

Example 1: $\{(c, c, c), (c, i, i), (c, i, i)\} \rightarrow \text{Outcome} = i$

¹¹ Landmore (2013: 90) says bluntly “all-inclusiveness is simply not feasible, and . . . representative democracy so far remains the only option for our mass societies . . .”

¹² Scott E. Page is the primary author on the appendix.

The first subset has 3 voters who make the correct choice and the other 2 subsets have exactly 1 voter who makes the correct choice. Thus with 5 voters making the correct choice the representative outcome is still incorrect.

Example 2: $\{(c, c, i), (c, c, i), (i, i, i)\} \rightarrow \text{Outcome} = c$

In both of the first 2 subsets exactly 2 of the voters make the correct choice, and in the third subset none do. Thus the partitions in Example 2 produce the correct outcome even though only 4 of the voters make the correct choice. These examples lead to Observation 5:

Observation 5. With 9 voters partitioned into 3 sets of size 3, there are distributions of votes such that the representative hierarchy produces the correct outcome when direct majority voting does not. There also exist distributions of votes in which the opposite holds: direct majority voting produces the correct outcome while representative hierarchy does not.

Neither form of aggregation dominates the other in the sense of always producing at least as good an outcome for any distribution of votes. What remains to be established is whether one institution is better *on average* given our assumptions regarding the individuals' probability of correctness. The following claim provides an answer to that question.

Claim 1. With 9 voters, direct majority rule has a strictly higher probability of producing the correct outcome than a representative hierarchy in which the voters are partitioned into 3 sets of size 3.

Proof. As before we use c to denote the correct outcome and i for the incorrect outcome. From above, we know that if 6 or more, or 3 or fewer, voters vote c then we get the same outcome for direct majority and representative hierarchy voting. Let $R(4, c)$ denote the set of all realizations in which exactly 4 of the voters vote c and the outcome of a representative hierarchy is c . Designate the sets $R(4, i)$, $R(5, c)$ and $R(5, i)$ similarly. The sets $R(4, i)$ and $R(5, c)$ produce the same outcome as direct majority; i.e. for any realization in $R(4, i)$, 4 of the voters vote c and the group outcome is i , the same as a direct majority vote with 4 c votes. The only cases in which a direct majority vote and representative democracy produce different outcomes are therefore $R(4, c)$ and $R(5, i)$. Within $R(4, c)$ a representative hierarchy produces the correct outcome and direct majority voting does not. Within $R(5, i)$ the opposite holds. To complete the proof, we rely on the following lemma regarding cardinality of the sets:

Lemma 1. $|R(5, i)| = |R(4, c)|$

The proof of the lemma is as follows: Consider any realization in $R(5, i)$. It will consist of 5 voters who vote c and 4 who vote i such that 2 of the 3 subsets in the partition have a majority of voters who vote i . The only way this can happen is if 2 of the subsets contain 2 voters who vote i , as in example 1:

Example 1: $\{(c, c, c), (c, i, i), (c, i, i)\} \rightarrow \text{Outcome} = i$

Given this example, we can construct a mirror realization in which every c signal becomes an i and every i signal becomes a c .

Mirror of Example 1: $\{(i, i, i), (i, c, c), (i, c, c)\} \rightarrow \text{Outcome} = c$

The lemma follows from two observations. First, any realization has a unique mirror realization. Second, the mirror realization to any realization in $R(5, i)$ lies in $R(4, c)$ and vice versa.

To complete the proof of Claim 1 we only need to show that any realization in $R(5, i)$ has a higher probability than its mirror realization in $R(4, c)$. This follows directly from the fact that individuals vote c with a probability $p > 0.5$. The probability of any particular realization in $R(5, i)$ – the cases in which a direct majority produces a correct result while a representative hierarchy does not – is

$p^5(1-p)^4$. The probability of the mirror realization in $R(4, c)$ is $p^4(1-p)^5$, which is always smaller than $p^5(1-p)^4$ when $p > 0.5$, thus completing the proof.

General Result

We can now state a general result given n voters partitioned into g (odd) subsets.

Claim 2. With n voters, a direct majority vote has a strictly higher probability of producing the correct outcome than a representative hierarchy in which the voters are partitioned into g (odd) subsets each of which has an odd (and not necessarily equal) number of voters.

Proof. Let $R(v, x)$ denote the set of all voting realizations in which v voters choose c and the representative hierarchy produces the outcome x . Then let D_C denote the set of all $R(v, x)$ s in which direct majority voting produces a correct outcome while a representative hierarchy does not, and let D_I denote the set of $R(v, x)$ s where the opposite holds.

By construction:

$$D_C = \{R(v, i) \& v > n/2\}$$

$$D_I = \{R(v, c) \& v < n/2\}$$

For each set of $R(v, i)$ s in D_C define its partners as $R(n-v, c)$. Note that partners create a one-to-one mapping between sets in D_C and sets in D_I . Suppose that there exists at least one realization in $R(v, i)$; i.e., that $R(v, i)$ is non-empty. As above, define the mirror realization by changing every c vote to an i vote and every i vote to a c . It follows that the mirror realization of a realization in $R(v, i)$ belongs to its partner set $R(n-v, c)$. Similarly, every realization in $R(n-v, c)$ has a unique mirror realization in $R(v, i)$. Therefore, there are the same number of realizations in $R(v, i)$ as there are in $R(n-v, c)$. Also following the logic from above, the probability of any given realization in $R(v, i)$ equals $p^v(1-p)^{n-v}$ which strictly exceeds the probability of its mirror realization with probability $p^{n-v}(1-p)^v$ as long as $v > n/2$. This implies that each set in D_C has a higher probability of occurring than its corresponding set in D_I , which completes the proof.

REFERENCES

- Allen, T. J. and Cohen, S. I. 1969. 'Information Flow in Research and Development Laboratories.' *Administrative Science Quarterly*, 14: 12–19.
- Anderson, C. and Brown, C. E. 2010. 'The Functions and Dysfunctions of Hierarchy.' *Research in Organizational Behavior*, 30: 55–89.
- Anderson, E. 2006. 'The Epistemology of Democracy.' *Episteme*, 3: 8–22.
- Austen-Smith, D. and Feddersen, T. J. 2005. 'Deliberation and Voting Rules.' In Austen-Smith D. and Duggan J. (eds) *Social Choice and Strategic Decisions: Essays in Honor of Jeffrey S. Banks*, pp. 269–316. Heidelberg: Springer.
- 2006. 'Deliberation, Preference Uncertainty, and Voting Rules.' *American Political Science Review*, 100: 209–17.
- Bednar, J. 2009. *The Robust Federation: Principles of Design*. New York, NY: Cambridge University Press.
- Beisbart, C. and Bovens, C. E. 2013. 'Minimizing the Threat of a Positive Majority Deficit in a Two-tier Voting System with Equipopulous Units.' *Public Choice*, 145: 75–94.
- Betz, G. 2013. *Debate Dynamics: How Controversy Improves our Beliefs*, Vol. 357. Heidelberg: Springer Science & Business Media.

- Boland, P. J. 1989. 'Majority Systems and the Condorcet Jury Theorem.' *Journal of the Royal Statistical Society. Series D (The Statistician)*, 38: 181–9.
- , Proschan, F. and Tong, Y. L. 1989. 'Modeling Dependence in Simple and Indirect Majority Systems.' *Journal of Applied Probability*, 26: 81–8.
- Cantimur, Y. 2016. 'When and why hierarchy steepness is related to team performance.' *European Journal of Work and Organizational Psychology*, 25: 658–73.
- Cohen, J. 1986. 'An Epistemic Conception of Democracy.' *Ethics*, 97: 26–38.
- Condorcet, M. 1785 [1995]. 'An Essay on the Application of Analysis to the Probability of Decisions Rendered by a Plurality of Votes (Fifth Part).' In I. McLean and A. B. Urken (transl. and ed.), *Classics of Social Choice*, pp. 91–112. Ann Arbor, MI: University of Michigan Press.
- Dahl, R. A. 1956. *A Preface to Democratic Theory*. Chicago, IL: University of Chicago Press.
- Dahl, R. A. 1983. *Dilemmas of Pluralist Democracy: Autonomy versus Control*. New Haven, CT: Yale University Press.
- Dietrich, F. and Spiekermann, K. 2013. 'Epistemic Democracy with Defensible Premises.' *Economics and Philosophy*, 29: 87–120.
- , List, C. and Bradley, R. 2016. 'Belief Revision Generalized: A Joint Characterization of Bayes' and Jeffrey's rules.' *Journal of Economic Theory*, 162: 352–71.
- Estlund, D. 1994. 'Opinion Leaders, Independence, and Condorcet's Jury Theorem.' *Theory and Decision*, 36: 31–162.
- 2008. *Democratic Authority: A Philosophical Framework*. Princeton, NJ: Princeton University Press.
- Fishkin, J. S. 1991. *Democracy and Deliberation: New Directions for Democratic Reform*. New Haven, CT: Yale University Press.
- 1995. *The Voice of the People: Public Opinion and Democracy*. New Haven, CT: Yale University Press.
- Fishkin, J. S. 2009. *When the People Speak: Deliberative Democracy and Public Consultation*. New York, NY: Oxford University Press.
- Fung, A. 2006. 'Democratizing the Policy Process.' In R. Goodin, M. Moran and M. Rien (eds), *The Oxford Handbook of Public Policy*, pp. 669–85. New York, NY: Oxford University Press.
- Gaus, G. 1997. 'Does Democracy Reveal the Voice of the People? Four Takes on Rousseau.' *Australasian Journal of Philosophy*, 75: 141–62.
- Goodin, R. E. and Spiekermann, K. 2012. 'Epistemic Aspects of Representative Government.' *European Political Science Review*, 4: 303–25.
- Grofman, B. and Feld, S. L. 1988. 'Rousseau's General Will: A Condorcetian Perspective.' *American Political Science Review*, 82: 567–76.
- Grönlund K., Bächtiger A. and Setälä M. (eds) 2013. *Deliberative Mini-Publics: Practices, Promises, Pitfalls*. New York, NY: Columbia University Press.
- Habermas, J. 1996. *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy* (Transl. T. Berger). Cambridge, MA: MIT Press.
- Hong, L. and Page, S. E. 2004. 'Groups of Diverse Problem Solvers can Outperform Groups of High-ability Problem Solvers.' *Proceedings of the National Academy of Sciences USA*, 101: 16385–9.
- Kohn, N. W., Paulus, P. B. and Choi, Y. 2011. 'Building on the Ideas of Others: An Examination of the Idea Combination Process.' *Journal of Experimental Social Psychology*, 47: 554–61.
- Kuehn, D. 2017. 'Diversity, Ability, and Democracy: A Note on Thompson's Challenge to Hong and Page.' *Critical Review*, 29: 72–87.
- Ladha, K. K. 1992. 'The Condorcet Jury Theorem, Free Speech, and Correlated Votes.' *American Journal of Political Science*, 36: 617–34.
- 1995. 'Information Pooling Through Majority-rule Voting: Condorcet's Jury Theorem with Correlated Votes.' *Journal of Economic Behavior and Organization*, 26: 353–75.
- Landa, D. and Meirowitz, A. 2009. 'Game Theory, Information, and Deliberative Democracy.' *American Journal of Political Science*, 53: 427–44.
- Landmore, H. 2013. *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton, NJ: Princeton University Press.
- and Elster J. (eds) 2012. *Collective Wisdom: Principles and Mechanisms*. New York, NY: Cambridge University Press.

- and Page, S. E. 2015. 'Deliberation and Disagreement: Problem Solving, Prediction, and Positive Dissensus.' *Philosophy, Politics, and Economics*, 14: 229–54.
- Levine, J. M. and Smith, E. R. 2013. 'Group Cognition: Collective Information Search and Distribution.' In D. E. Carlston (ed.), *The Oxford Handbook of Social Cognition*, pp. 616–33. New York, NY: Oxford University Press.
- and Tindale, R. S. 2015. 'Social Influence in Groups.' In M. Mikulincer, P. R. Shaver, J. F. Dovidio and J. A. Simpson (eds), *APA Handbook of Personality and Social Psychology, Volume 2: Group Processes*, pp. 3–34. Washington, DC: American Psychological Association.
- Lijphart, A. 1999. *Patterns of Democracy: Government Forms and Performance in Thirty-six Countries*. New Haven, CT: Yale University Press.
- List, C. 2005. 'Group Knowledge and Group Rationality: A Group Aggregation Perspective.' *Episteme*, 2: 25–38.
- 2018. 'Democratic Deliberation and Social Choice: A Review.' In A. Bächtiger, J. S. Dryzek, J. Mansbridge and M. E. Warren (eds), *The Oxford Handbook of Deliberative Democracy*. New York, NY: Oxford University Press.
- List, C. and Goodin, R. E. 2001. 'Epistemic Democracy: Generalizing the Condorcet Jury Theorem.' *Journal of Political Philosophy*, 9: 277–306.
- Mathis, J. 2011. 'Deliberation with Evidence.' *American Political Science Review*, 105: 516–29.
- Miller, N. R. 1986. 'Information, Electorates, and Democracy: Some Extensions and Interpretations of the Condorcet Jury Theorem.' In B. Grofman & G. Owens (eds), *Information Pooling and Group Decision Making: Proceedings of the Second University of California, Irvine Conference on Political Economy*, pp. 173–192. Greenwich: JAI Press.
- Niemeyer, S. 2011. 'The Emancipatory Effect of Deliberation: Empirical Lessons from Mini-publics.' *Politics & Society*, 39: 103–40.
- Page, S. E. 2007. *The Difference*. Princeton, NJ: Princeton University Press.
- 2011. *Diversity and Complexity*. Princeton, NJ: Princeton University Press.
- Perote-Peña, J. and Piggins, A. 2015. 'A Model of Deliberative and Aggregative Democracy.' *Economics and Philosophy*, 31: 93–121.
- Pitkin, H. F. 1967. *The Concept of Representation*. Berkeley, CA: University of California Press.
- Rittel, H. W. J. 1965. 'Hierarchy or Team? Considerations on the Organization of the R&D Cooperatives.' In R. A. Tybur (ed.), *Economics of Research and Development*, pp. 174–218. Columbus, OH: Ohio State University Press.
- Rousseau, J.-J. 1762 [1978]. *The Social Contract*. (Transl. J. Masters and R. Masters). New York, NY: St. Martins.
- Schwartzberg, M. 2015. 'Epistemic Democracy and its Challenges.' *Annual Review of Political Science*, 18: 187–203.
- Singer, D. Forthcoming. 'Diversity, Not Randomness, Trumps Ability.' *Philosophy of Science*.
- Stasser, G. and Dietz-Uhler, B. 2001. 'Collective Choice, Judgment, and Problem Solving.' In M. A. Hogg and R. S. Tindale (eds), *Blackwell Handbook of Social Psychology: Group Processes*, pp. 31–55. Oxford: Blackwell.
- Thompson, A. 2014. 'Does Diversity Trump Ability? An Example of the Misuse of Mathematics in the Social Sciences.' *Notices of the American Mathematical Society*, 61: 1024–30.
- Tindale, R. S., Nadler, J., Krebel, A. and Davis, J. H. 2001. 'Procedural Mechanisms and Jury Behavior.' In M. A. Hogg & R. S. Tindale (eds), *Blackwell Handbook of Social Psychology: Group Processes*, pp. 574–602. Oxford: Blackwell.
- Urbanian, N. and Warren, M. E. 2008. 'The Concept of Representation in Contemporary Democratic Theory.' *Annual Review of Political Science*, 11: 387–412.
- Warren M. E. and Pearse H. (eds) 2008. *Designing Deliberative Democracy: The British Columbia Citizens' Assembly*. New York, NY: Cambridge University Press.
- Witte E. and Davis J. H. (eds) 2014. *Understanding Group Behavior: Consensual Action by Small Groups*, Vol. 1. New York, NY: Psychology Press.

This work was produced by the COMPUTATIONAL SOCIAL PHILOSOPHY LAB (CSPL), an interdisciplinary and multi-institutional group of researchers that investigates

questions of social epistemology, philosophy of science, and political philosophy using traditional philosophical methods and agent-based computer models with an eye toward empirical results.

WILLIAM J. BERGER is a fellow with the Philosophy, Politics, and Economics program at the University of Pennsylvania. He works on topics of polarization, inequality, and trust that lie at the intersection of empirical and normative political science.

AARON BRAMSON is a research scientist in the Lab for Symbolic Cognitive Development at the Riken Brain Science Institute in Japan and the Department of General Economics at the University of Gent in Belgium as well as an adjunct professor in the Dept of Software and Information Systems of UNCC in the US. He studies the foundations of complexity and specializes in the development of new methodologies and conceptual frameworks for the study of complexity itself and the systems that evoke it.

PATRICK GRIM is Distinguished Teaching Professor of Philosophy Emeritus at Stony Brook University and Philosopher in Residence with the Center for Study of Complex Systems at the University of Michigan. He is author of *The Incomplete Universe* and co-author of *Beyond Sets, Reflexivity: From Paradox to Consciousness*, and *The Philosophical Computer*. He is editor of *Mind and Consciousness: 5 Questions*, and founding co-editor of over 35 volumes of the *Philosopher's Annual*. He publishes widely in scholarly journals across several disciplines, and is incoming editor of the *American Philosophical Quarterly*. Contact Email: patrick.grim@stonybrook.edu

JIIN JUNG is a doctoral candidate in Social Psychology at the Claremont Graduate University. Her research focuses on uncertainty reduction as an epistemic motive of social identity and self-categorization processes. She also examines computational models of social influence processes and uncertainty-related determinants of depersonalization and projection.

SCOTT E. PAGE is the Leonid Hurwicz Collegiate Professor of Complex Systems, Political Science, and Economics at the University of Michigan and an external faculty member of the Santa Fe Institute. A fellow of the American Academy of Arts and Sciences and a Guggenheim Fellow, he studies the functional contributions of diversity in complex systems, mechanism and institutional design, and the application of formal models. His most recent book, *The Model Thinker* (2018) describes the value of using multiple models.

DANIEL J. SINGER is an Assistant Professor of Philosophy at the University of Pennsylvania. His research is primarily at the intersections of epistemology, ethics, and social philosophy.
