

Cognitive Diversity, Binary Decisions, and Epistemic Democracy

John A. Weymark

Abstract

In *Democratic Reason*, Helene Landemore has built a case for the epistemic virtues of inclusive deliberative democracy based on the cognitive diversity of the group engaged in making collective decisions. She supports her thesis by appealing to the Diversity Trumps Ability Theorem of Lu Hong and Scott Page. In practice, deliberative assemblies often restrict attention to situations with only two options. In this paper, it is shown that it is not possible to satisfy the assumptions of the Diversity Trumps Ability Theorem when decisions are binary. The relevance of this theorem for democratic decision-making in non-binary situations is also considered.

Working Paper: 5-2014

Research Concentration: Elections and Electoral Rules

Cognitive Diversity, Binary Decisions, and Epistemic Democracy

John A. Weymark

*Departments of Economics and Philosophy, Vanderbilt University,
PMB 351819, 2301 Vanderbilt Place, Nashville, TN 37235-1819, USA*

john.weymark@vanderbilt.edu

July 2014

Revised, August 2014

Abstract. In *Democratic Reason*, H el ene Landemore has built a case for the epistemic virtues of inclusive deliberative democracy based on the cognitive diversity of the group engaged in making collective decisions. She supports her thesis by appealing to the Diversity Trumps Ability Theorem of Lu Hong and Scott Page. In practice, deliberative assemblies often restrict attention to situations with only two options. In this paper, it is shown that it is not possible to satisfy the assumptions of the Diversity Trumps Ability Theorem when decisions are binary. The relevance of this theorem for democratic decision-making in non-binary situations is also considered.

I have benefited from correspondence and discussions with H el ene Landemore, Scott Page, and Robert Talisse, as well as from the comments received when this paper was presented at the 2014 Meeting of the Society for Social Choice and Welfare.

1. Introduction

Hélène Landemore (2013) has recently built a case for the epistemic virtues of inclusive deliberative democracy based on the cognitive diversity of the group engaged in making collective decisions. By *inclusive*, Landemore means that all members of the relevant group take part in the deliberation. *Cognitive diversity* refers to “the variety of mental tools that human beings use to solve problems or make predictions in the world”—their mental toolkits (p. 89). Landemore supports her thesis by appealing to the Diversity Trumps Ability Theorem of Lu Hong and Scott Page, first presented formally in Hong and Page (2004) and then informally in Page (2007). This theorem shows that, in some circumstances, a sufficiently large number of individuals chosen randomly from a cognitively diverse group can be expected to be better at identifying the best outcome than a group of experts with the same number of members. In other words, cognitive diversity brings with it collective wisdom.¹

The Diversity Trumps Ability Theorem applies whatever the number of options available to the deliberative group. In practice, deliberative assemblies often restrict attention to situations with only two options—a *binary decision*.² In this paper, I show that it is not possible to satisfy the assumptions of the Diversity Trumps Ability Theorem when there are only two options being considered. Hence, this theorem cannot provide support for the epistemic virtues of inclusive democratic deliberation for binary decisions. In the concluding section, I provide some remarks on the relevance of the Diversity Trumps Ability Theorem for democratic decision-making in non-binary situations.

Hong and Page’s Diversity Trumps Ability Theorem is concerned with a group of individuals who share the same standard for the goodness of decisions and whose goal is to determine the best outcome according to this standard given their limited cognitive abilities. The group is thus engaged in a cooperative problem-solving exercise.³ The theorem applies to a cognitively

¹Many examples of this phenomenon are provided in Page (2007) and Surowiecki (2004).

²In her discussion of majority rule, Landemore (2013, pp. 149–150) argues that there are good reasons, both empirical and theoretical, for restricting attention to situations in which voting is over two options.

³Democratic assemblies also engage in other sorts of activities, such as aggregating preferences and making predictions about the consequences of adopting some policy. The latter is an epistemic activity, whereas the former is not. On the differences between these

diverse group of reasonably competent individuals tackling a difficult problem for which there is a best, but *a priori* unknowable, way to search for a solution.

The statement of the Diversity Trumps Ability Theorem in Landemore (2013) is informal and somewhat inaccurate.⁴ Furthermore, some of the concerns about Landemore’s application of this theorem to democratic politics raised by Gunn (2014) and Quirk (2014) are misguided because they are based on an incorrect understanding of what the theorem actually establishes. The original statement of the Diversity Trumps Ability Theorem in Hong and Page (2004) is rather terse and technical, and therefore inaccessible to many scholars interested in its implications for democratic theory. There is therefore a need for an exposition of this theorem that does not sacrifice precision for accessibility. One of the objectives of this paper is to fill this need. Moreover, to show that the Diversity Trumps Ability Theorem is vacuous when there are only two options also requires a precise statement of the theorem and its assumptions.

Cohen (1986, p. 34) has famously characterized an epistemic approach to democracy by three desiderata. First, there exist correct decisions according to some procedure-independent standard of correctness. For example, for Cohen this standard is provided by an account of justice or the common good, whereas for Estlund (2008) the standard is moral truth. Second, democratic decision-making is about individuals expressing their beliefs about what the correct decision is, not about their values and personal preferences. That is, democratic decision-making is inherently cognitive. Third, the decision-making process is one of adjusting individual beliefs in response to the beliefs about the correct decision provided by others. There is thus a deliberative aspect to epistemic democracy.

For an epistemic democrat, one (and for some epistemic democrats, the only) basis on which the authority and legitimacy of democratic procedures rest is their ability to make correct decisions, that is, on their ability to “track the truth”. The authority and legitimacy of a particular democratic institution is enhanced to the extent that it generally tracks the truth better than alternative forms of collective decision-making.⁵ Landemore’s application

kinds of activities and their implications for democratic decision-making, see Landemore (2013) and Landemore and Page (2014).

⁴The same is true about the presentation of this result in Anderson (2006).

⁵There are many different versions of epistemic democracy. For an introduction and critical analysis of the main contenders, see Peter (2009).

of the Diversity Trumps Ability Theorem to democratic decision-making belongs to this tradition. She is therefore considering an alternative mechanism for tracking the truth than the one provided by the Condorcet Jury Theorem.⁶ While Landemore (2013) focuses on epistemic success, as she makes clear in Landemore (2014), the authority and legitimacy of democratic institutions also depend on non-epistemic factors, such as the expressive or intrinsic value of the procedures used to make collective decisions.

One criterion that has been used to judge how well a collective decision-making procedure tracks the truth is that it is more reliable in identifying the correct decision than a random procedure (Estlund, 2008, p. 160). The Democracy Trumps Ability Theorem shows that the democratic procedures to which it applies satisfy a much more demanding criterion: they do better than a group of experts. Thus, to the extent that this theorem does provide support for deliberative democratic decision-making by a cognitively diverse set of individuals in non-binary situations, it passes a stringent test for success.

The plan of the rest of this paper is as follows. In Section 2, the model of individual decision-making employed in the Diversity Trumps Ability Theorem is introduced. Section 3 discusses how the kinds of problems that this theorem applies to is formalized in terms of the procedures individuals use to search for good decisions. How collective decisions are made is described in Section 4. The theorem itself is presented in Section 5. The impossibility of satisfying all of the assumptions of the Diversity Trumps Ability Theorem is established in Section 6. Finally, in Section 7, some comments on the relevance of this theorem for democratic decision-making in non-binary situations are provided.

2. Individual Decision-Making

The collective decision problem that a group of individuals faces is one of identifying the best alternative from some set X of possible choices. Here, X assumed to be finite. Of course, X must contain at least two alternatives, otherwise there is no choice to be made. For example, if the question

⁶Relatively non-technical introductions to the Condorcet Jury Theorem are provided in Estlund (2008), List and Pettit (2011), and Landemore (2013). There is considerable controversy about the relevance of the Condorcet Jury Theorem for democratic decision-making. Some of the sceptic's concerns may be found in Estlund (2008), Peter (2009), and Landemore (2013).

is whether to mandate that automobiles have anti-theft alarms installed, whether to ban smoking in bars, or to determine whether a defendant in a trial is guilty, then there are two choices: {Yes, No}. Nonbinary choices are also possible. Page (2007, pp. 139–143) considers an example in which the problem is one of determining which of 1,000 automobiles on a car dealer’s lot gets the best gas mileage. To help motivate her discussion of the Diversity Trumps Ability Theorem, Landemore (2013, pp. 99–100) contemplates the problem of determining which city in France the government should locate an experimental programme in. As a final example, in a Scottish jury trial, there are three options: {Guilty, Not Guilty, Not Proven}.

It is assumed that there is some procedure-independent standard for determining the goodness of the alternatives that is shared by everyone. This measure of goodness is modeled by a *value function* $V: X \rightarrow \mathbb{R}$ that assigns a value $V(x)$ to each alternative x in X , with larger values assigned to better alternatives. It is supposed that there is a uniquely best alternative x^* in X . In the car example, the value of a car is its actual gas mileage. In the binary examples, the better decision could be assigned a value of 1 and the other decision a value of 0. Thus, in a jury problem with guilt and innocence as the only options, if the defendant is actually guilty, the value of convicting him would be 1, but if he is innocent, the value of conviction would be 0. If, as in Estlund (2008), the problem is one of determining moral truths, then the moral truth is assigned the highest value.

The individual decision makers have limited cognitive abilities. This manifests itself in a number of ways. Individuals may conceptualize the alternatives in different ways and focus on particular features of the alternatives so as to render the evaluation of alternatives less cognitively demanding. In the terminology of Page (2007), they have different *perspectives* (how alternatives are conceptualized) and *interpretations* (how the alternatives are grouped into coarser categories). In addition, because of their limited cognitive abilities, individuals may not be able to evaluate all of the alternatives. Even when there are only two alternatives, it may be sufficiently demanding to determine the value of an alternative that only one of them is chosen for evaluation. How someone structures the way he searches for a good decision is described by a *heuristic*. This heuristic may take account of the time, effort, and resources needed to keep searching, not just the cognitive difficulty of doing so.

For the purposes of the Diversity Trumps Ability Theorem, the fine structure of an individual’s perspective, interpretation, and heuristic do not mat-

ter. What matters is the alternative he initiates his search with and the alternative his search for a good decision terminates at. Hong and Page (2004) model this procedure by specifying a starting rule and a search rule. A *starting rule* ν is a probability distribution on X with $\nu(x)$ denoting the probability that the search process starts with alternative x . It is assumed that everybody uses the same starting rule ν and that this probability distribution has full support.⁷ A *search rule* $\phi: X \rightarrow X$ specifies what alternative $\phi(x)$ to recommend for each starting point $x \in X$. Many starting points could terminate with the same recommendation. The collection of alternatives at which the search procedure could stop is $\phi(X) = \{\bar{x} \in X: \bar{x} = \phi(x) \text{ for some } x \in X\}$. This is the set of *local optima* for an individual with search rule ϕ .

As an illustration, consider the gas mileage example. Suppose that an individual believes (perhaps incorrectly) that gas mileage is strongly correlated with wheel base length. He first ranks the cars by the lengths of their wheel bases. If there are only a few makes and models of car on the lot, this is a relatively easy task. He then randomly picks a car to inspect first. After an internet search, he discovers the gas mileage for this car. Next, he investigates a car with the next highest wheel base (if there is one). If it has better gas mileage, he repeats this procedure until either doing so identifies a car with lower gas mileage (in which case he returns to the previous car) or the car has the longest wheel base on the lot. If the second car investigated has a lower gas mileage than the first car considered or if the first car has the longest wheel base on the lot, this procedure is run in reverse by evaluating cars with smaller wheel bases. The car at which this search terminates is a local optimum for this decision maker. If gas mileage is not perfectly correlated with wheel base, had a different car been used to initiate the search, this heuristic may identify a different car—one of this individual’s other local optima. If a different individual employs a different perspective and interpretation, say by grouping cars by their weight, but uses the same search heuristic reinterpreted in terms of weight instead of length of wheel base, he would have a different set of local optima.

Someone might start his search near the best alternative x^* and find it with little effort. This fortuitous outcome is a matter of luck, not expertise. In the Diversity Trumps Ability Theorem, an individual’s expertise—his in-

⁷The *support* of a probability distribution is the set of alternatives that have positive probability. Thus, to say that ν has full support means that the search could start with any of the alternatives in X .

dividual performance—is an *ex ante* measure of how good a job the search and starting rules he uses do in identifying good alternatives. Formally, the *individual performance* of someone with search rule ϕ and starting rule ν is given by the expected value $E(V; \phi, \nu)$ of the outcome that is identified using these rules, where

$$E(V; \phi, \nu) = \sum_{x \in X} V(\phi(x))\nu(x). \quad (1)$$

In this formula, first, for each starting point x , the value of the outcome that is reached from it is weighted by the probability of starting the search there and then these probability-weighted values for all possible starting points are summed. An individual who uses the rules (ϕ, ν) *exhibits more expertise than* an individual who uses the rules (ϕ', ν') if and only if

$$E(V; \phi, \nu) > E(V; \phi', \nu'). \quad (2)$$

Thus, the experts are those individuals whose search and starting rules result in the highest levels of individual performance. As noted earlier, here it is assumed that everybody uses the same starting rule. When this is the case, an individual’s expertise is being judged by the effectiveness of his search rule.

3. Admissible Search Rules

The Diversity Trumps Ability Theorem does not apply to all collective decisions. Rather, it is concerned with difficult problems for which the individual decision makers are reasonably competent and have a diversity of approaches to solving such problems and for which there is a uniquely best search rule. These four requirements—difficulty, competence, diversity, and the existence of a uniquely best search rule—are formalized by placing restrictions on the set of *admissible search rules* Φ that are used by the individuals making the collective decision. These restrictions are formalized axiomatically.

Assumption 1 (Competence). For all $\phi \in \Phi$ and all $x \in X$,

1. $V(\phi(x)) \geq V(x)$, and
2. $\phi(\phi(x)) = \phi(x)$.

The first part of this assumption says that no matter which alternative x the search process starts with, it does not terminate at an alternative $\phi(x)$ that is worse than x . The second part says that if an individual reaches one of his local optima, then he cannot restart the search process from there and come up with a different alternative. If the search rule ϕ can identify a better alternative at $\phi(x)$, then this individual should not have terminated the search there. This is not to say that there are not better alternatives than $\phi(x)$, only that the search rule ϕ cannot find them if the search ever gets to $\phi(x)$.⁸ Violating either of the conditions in Assumption 1 is a sign that this decision maker is not very competent. When combined with the assumption that the search can start with any alternative in X , an important implication of Assumption 1 is that x^* is a local optimum of any search rule that satisfies it.

Assumption 2 (Difficulty). For all $\phi \in \Phi$, there exists an $x \in X$ such that $\phi(x) \neq x^*$.

This assumption captures the idea that the problem of identifying the best alternative is sufficiently difficult that nobody on his own is always going to be able to solve it. When a problem is difficult, it is natural to suppose that there are many alternatives from which the search could be initiated that terminate before identifying the best outcome. However, it is sufficient for the Diversity Trumps Ability Theorem that there is at least one such starting point.

Assumption 3 (Diversity). For all $x \in X \setminus \{x^*\}$, there exists a $\phi \in \Phi$ such that $\phi(x) \neq x$.

Every individual uses a search rule in Φ . Assumption 3 ensures that no alternative other than x^* is a local optimum for all of these search rules. If someone terminates his search at $x \neq x^*$, there is another search rule that recognizes that better alternatives exist. In other words, the search rules in Φ encapsulate a diversity of approaches to finding the best alternative. The assumption that any non-optimal alternative can be improved on by some search rule in Φ does not imply that any particular group of decision makers uses it. Indeed, if a group is sufficient homogeneous in its choice of search

⁸Assumption 1 can be thought of as an analogue of how a hill-climbing procedure is used to find a local maximum of a real-valued function—never go downhill and always stop when a local peak is encountered.

	\bar{x}	\hat{x}	x^*
ϕ^1	\bar{x}	\hat{x}	x^*
ϕ^2	\bar{x}	x^*	x^*
ϕ^3	x^*	\hat{x}	x^*

Table 1: An admissible class of search rules for a tertiary decision.

rules, which is the case with a large group of experts, they may well get stuck at one of their local optima that differs from x^* because nobody in this group is using a search rule that recognizes that further improvements are possible. In contrast, with a group that utilizes a diverse set of search rules, getting stuck at a non-optimal alternative is less likely to happen.

Assumption 4 (Unique Best Search Rule). $\arg \max\{E(V; \phi, \nu) : \phi \in \Phi\}$ is unique.⁹

This assumption says that there is exactly one search rule that outperforms all the others. An individual who uses this rule exhibits the greatest expertise possible. Individual performance depends on the starting rule ν , not just the search rule ϕ . However, here, everybody uses the same fixed starting rule, so Assumption 4 can be thought of as being an assumption just about the set of admissible search rules. Note that this is the only one of the four assumptions that takes account of the starting rule.

The set of search rules in the following example of a three-alternative collective decision problem satisfies all four assumptions.

Example. The set of alternatives is $X = \{\bar{x}, \hat{x}, x^*\}$ with $V(\bar{x}) < V(\hat{x}) < V(x^*)$. Thus, x^* is the best alternative, \hat{x} is second best, and \bar{x} is worst. All that is assumed about the starting rule ν is that ν has full support with $\nu(\bar{x}) = \nu(\hat{x})$. Thus, it is equally likely to start the search with either \bar{x} or \hat{x} and there is some chance that it starts at x^* . The set of admissible search rules is $\Phi = \{\phi^1, \phi^2, \phi^3\}$ defined as in Table 1. In this table, the entries in a column indicate what local optimum is reached by each of the three search rules when the alternative that heads the column is used as the starting point.

⁹*Arg max* means “the argument that maximizes”. Here, it refers to the search rules that result in the best individual performance.

It is readily verified that Φ satisfies the first three assumptions. Note in particular that the search terminates at x^* if it is the starting point, as required by Assumption 1. Because of this restriction, the ordering of search rules in terms of their performance is completely determined by the entries in the first two columns of the table. The first search rule terminates where it starts. With the other two rules, the one starting point that is not a local optimum terminates at a better alternative, so either of these rules perform better than ϕ^1 . Because $\nu(\bar{x}) = \nu(\hat{x})$, ϕ^2 and ϕ^3 are equally likely to identify x^* . The probability that ϕ^2 terminates at \bar{x} is the same as the probability that ϕ^3 terminates at \hat{x} . Because \hat{x} is the better of these two alternatives, the performance of ϕ^3 is better than that of ϕ^2 , which makes it the uniquely best search rule in Φ .

4. Collective Decision-Making

I now consider how a group of individuals makes its decision. What they decide as a group depends on the search rules that they employ and on how they utilize the information obtained from their searches.¹⁰

Suppose that there are n individuals, denoted by $N = \{1, \dots, n\}$. The i th of these individuals uses search rule $\phi_i \in \Phi$. The *profile* of search rules is $\mathbf{P} = (\phi_1, \dots, \phi_n)$. The set of local optima for individual i is $\phi_i(X)$. The set of *common local optima* for the profile \mathbf{P} is

$$C(\mathbf{P}) = \bigcap_{i \in N} \phi_i(X). \quad (3)$$

The alternatives in $C(\mathbf{P})$ consist of every alternative that all of the individuals in the group N agree are local optima. This set is nonempty because, by Assumption 1, x^* is a local optimum for every admissible search rule.

The choice made by a group with profile \mathbf{P} is denoted by $D(\mathbf{P})$. In the Diversity Trumps Ability Theorem, it is assumed that the group decision $D(\mathbf{P})$ is one of the alternatives in the set of common local optima $C(\mathbf{P})$. Furthermore, it is assumed that if there is more than one common local optimum, then the choice is made by employing a profile-dependent probability distribution $\eta_{\mathbf{P}}$ on X whose support is $C(\mathbf{P})$. Thus, it is possible for any of the common local optima to be chosen, but there is no chance that any other

¹⁰Strictly speaking, their decision also depends on the starting rules used, but because everybody uses the same fixed starting rule, this dependence does not need to be made explicit.

alternative is.¹¹

The assumption that $D(\mathbf{P})$ is chosen probabilistically from among the set of common local optima is consistent with a variety of more fully specified collective choice procedures. This assumption is a formalization of the idea that the individuals share a common interest in identifying the best alternative that they can given their cognitive limitations and that they will cooperate in order to do so, but that there are random elements in the decision-making process (such as the choice of where to start the search from). In particular, this way of modeling the collective choice problem can be interpreted as capturing some essential features of a deliberative democracy, and it is this interpretation that Landemore (2013) draws on.

One, but not the only, way that this model of collective decision-making can be interpreted in deliberative democratic terms builds on the sequential decision-making process considered by Hong and Page (2004) and Page (2007).¹² Imagine arranging the individuals clockwise in a circle preserving the order of their “names” (1, 2, etc.). Starting with individual 1 and moving clockwise, each of the individuals is encountered in turn. When individual n is reached, by continuing to move clockwise, the next person in line is individual 1. Individual 1 initiates the deliberative process. He uses the starting rule ν to determine which alternative x to begin his evaluation of the alternatives with. After completing his evaluation, he announces that $\phi_1(x)$ is the best alternative that he has found, what its value is, and an account of how he came to this conclusion. Then individual 2 determines whether he can improve on the first person’s finding by using his search rule starting at $\phi_1(x)$. This process is repeated sequentially by moving clockwise around the circle, as described above. This process terminates when a complete circuit of the circle has been made in which nobody has suggested an improvement. The alternative that is arrived at in this way is a common local optimum. Had individual 1 started his search with a different alternative, this sequential procedure could terminate in a different common local optimum. With this

¹¹Hong and Page (2004) only explicitly assume that the support of $\eta_{\mathbf{P}}$ is contained in $C(\mathbf{P})$. This assumption is compatible with choosing x^* for certain no matter what the profile is, in which case any group, diverse or not, has the maximal group performance. While it is sufficient for their theorem to assume that it is not always the case that x^* is chosen when there are multiple local optima, the stronger assumption used here is simpler to state and also quite natural to make.

¹²For an alternative way of modeling this sequential search procedure, see LiCalzi and Surucu (2012).

procedure, the probability distribution $\eta_{\mathbf{P}}$ used to select among the common local optima is determined by the starting rule ν and the profile of search rules \mathbf{P} . Admittedly, this is a highly stylized model of deliberation, but it captures some of its essential features—making proposals, giving reasons for them, responding in like manner to proposals made by other individuals, and sharing an interest in finding a mutually acceptable decision.¹³

The set of common local optima depends on which of the admissible search rules are used by the individuals. For the set of admissible search rules in the Example, suppose that the profile \mathbf{P} only contains individuals who use the search rules ϕ^1 and ϕ^2 . Then, $C(\mathbf{P}) = \{\bar{x}, x^*\}$ and either one of these alternatives is chosen with positive probability. If, instead, \mathbf{P} contains all three of the search rules in Φ , then $C(\mathbf{P}) = \{x^*\}$ and x^* (the best alternative) is chosen for sure.

The fact that x^* is the only common local optimum in the Example if every admissible search rule is used by some member of the group does not depend on the particulars of this example. By Assumption 1, x^* is a common local optimum for any profile of search rules. The diversity assumption (Assumption 3) ensures that no other alternative is a common local optimum if *all* of the search rules in Φ are used by the group. These observations play a fundamental role in the proof of the Diversity Trumps Ability Theorem.

The *joint performance* of a group of individuals with the profile \mathbf{P} is

$$\sum_{x \in X} V(x) \eta_{\mathbf{P}}(x). \quad (4)$$

That is, the group’s performance is measured by the expected value of the alternative it chooses.

5. The Diversity Trumps Ability Theorem

For the kind of collective choice problem considered here, the Diversity Trumps Ability Theorem established in Hong and Page (2004) shows that when selecting a large group of individuals from some larger group to make the collective decision, the joint performance of a randomly chosen set of individuals exceeds that of a group of the same size consisting of the best individual performers in the larger group. In other words, from an *ex ante*

¹³It is also possible for the individuals to simultaneously announce local optima and then have each of them check to see if he can improve on anybody else’s local optimum.

perspective, a diverse set of individuals does a better job of identifying the best alternative than a group of experts.

In the Diversity Trumps Ability Theorem, it is supposed that everybody uses a common starting rule ν with full support and that each of their search rules is chosen from a set of admissible search rules Φ that satisfies Assumptions 1–4. Furthermore, it is also supposed that any group chooses one of its common local optima, as described in the preceding section. Because individuals are only distinguished by the search rules they employ, choosing a group of individuals amounts to choosing what their search rules are. The group of experts is chosen by first picking a group of n individuals, each of whose search rules is independently chosen from Φ using a probability distribution μ on Φ that assigns a positive probability to each admissible search rule. The assumption that μ has full support ensures that when n is sufficiently large, the diversity of search rules found in Φ is preserved by the rules actually used by this group. From this group of n individuals, the n_1 best performers according to the individual performance measure $E(V; \phi, \nu)$ are chosen, where $n_1 < n$. This group is compared to another group of n_1 individuals whose search rules are independently chosen from Φ using the same probability distribution μ that is used to generate the set of individuals from whom the experts are chosen. Provided that n and n_1 are sufficiently large, the Diversity Trumps Ability Theorem shows that the randomly chosen group outperforms the group of experts.¹⁴ A precise formal statement of this result is provided by Theorem 1.

Theorem 1 (The Diversity Trumps Ability Theorem). *Suppose that X is finite and contains at least two alternatives, ν has full support, and Φ satisfies Assumptions 1–4. Furthermore, suppose that the choice $D(\mathbf{P})$ made by a group with profile \mathbf{P} is determined by a probability distribution $\eta_{\mathbf{P}}$ on X whose support is $C(\mathbf{P})$. Let μ be a probability distribution on Φ with full support. Then, there exist positive integers n and n_1 with $n_1 < n$ such that with probability one, the joint performance of n_1 individuals whose search rules are independently drawn from Φ according to μ exceeds the joint performance of the n_1 best decision makers among a group of n individuals whose search rules are independently drawn from Φ according to μ .¹⁵*

¹⁴If these groups are large, but not large enough for the conclusion of the theorem to hold, *ex ante* the randomly chosen group is expected to outperform the group of experts, but it need not do so with probability one.

¹⁵The theorem extends to non-finite sets of alternatives provided that it is possible to

There is simple intuition for this theorem. When the group chosen at random is sufficiently large, each search rule in Φ will be used by some individual. But then, as has been observed in the preceding section, x^* is the only common optimum for this group. Thus, the randomly chosen group correctly identifies the best possible alternative if it is sufficiently large. By Assumption 4, there is a uniquely best search rule in Φ . If the group of experts is sufficiently large, they will all use this rule. However, by Assumption 2, none of the admissible search rules finds x^* from every possible starting point. Thus, the experts could get stuck at one of the local optima for their common search rule different from x^* because of bad luck in the choice of starting point. The randomly chosen group never gets stuck in this way because there is always someone in the group who recognizes that any alternative other than x^* can be improved on. The diversity of their search rules trumps the fallible rule used by the experts.

The Example can be used to illustrate these observations. In this example, ϕ^3 is the best search rule. A sufficiently large group of best performers chosen as in the statement of the Diversity Trumps Ability Theorem all use this rule. As a group, they choose either \hat{x} or x^* , each with positive probability. So, there is some chance that they will make the wrong decision. If the same number of individuals is chosen randomly, then with probability one they will use every search rule in Φ , and so their only common local optimum is x^* . Thus, with this group, the correct decision is reached for sure.

It might appear that the Diversity Trumps Ability Theorem is making cognitive demands on the individual decision makers that are unachievable in practice. Each individual needs to determine the exact value of each alternative that it evaluates and then publicly announce the value of the local optimum that it identifies, or so it seems. In fact, this is not the case. For a group to make a collective decision in the way described above, it is sufficient for each individual to be able to recognize when one alternative has more value than another, but he does not need to know by how much. An individual terminates his search when he cannot find any better alternative given his limited ability to search for and evaluate alternatives. It is not necessary for him to announce this alternative's value, only that it is a local optimum for him. This information is sufficient for anybody else to have good reason to investigate whether his own search rule can improve on this choice. Moreover, even if he has announced the value of his local optimum,

count the number of local optima. See Page (2007).

nobody else need accept this announcement at face value without evaluating it himself. Thus, the information provided by others helps overcome one's own cognitive limits by suggesting further alternatives to consider.¹⁶

What the Diversity Trumps Ability Theorem tells us is that whoever the experts turn out to be, they will not outperform a randomly chosen set of decision makers. To prove this result, the actual values of the alternatives are needed to rank individuals in terms of performance and to compare the group performance of a random set of decision makers with that of a group of experts. However, once the epistemic superiority of the randomly chosen group has been established, there is no further need for these values because the random selection of the group members makes no use of these values and because there is no need to know who the experts are. In contrast, if an epistocracy had been assigned the task of making collective decisions, it would then be necessary to identify who the experts actually are, and that depends on the numerical values of the alternatives.

In fact, values do not even need to be measurable on an absolute scale in order to make comparative statements about performance. If values are subjected to an increasing affine transform, then the rankings of individuals by their performance as measured by (1) and groups by their performance as measured by (4) are unchanged.¹⁷ In the special case in which there are only two alternatives in X , it is only necessary to know which of them is better in order to compare performance because the two parameters that define an affine transform are sufficient to transform any pair of values into any other pair in an order-preserving way.

¹⁶Landemore and Page (2014) recognize that it is the relative rankings of alternatives that matter to the decision makers, not the alternatives' absolute values. They also say that they "implicitly assume the existence of an oracle, namely a machine, person, or internal intuition that can reveal the correct ranking of any proposed solutions." In other words, once anyone considers a pair of alternatives, which of them is better according to the standard used to determine relative value is apparent. This assumption is unnecessary. What matters is that when someone says that an alternative is a local optimum for him, this announcement induces someone else to consider this alternative. Its (relative) value may take the latter individual some effort to determine or verify.

¹⁷An *increasing affine transform* is a function of the form $f(t) = \alpha + \beta t$, where $\beta > 0$ and t is any real number. For example, temperature in the Celsius scale is converted to Fahrenheit using the affine transform in which $\alpha = 32$ and $\beta = 9/5$.

	\bar{x}	x^*
ϕ^a	\bar{x}	x^*
ϕ^b	x^*	\bar{x}
ϕ^c	\bar{x}	\bar{x}
ϕ^d	x^*	x^*

Table 2: All possible search rules for a binary decision.

6. Binary Decisions

In practice, deliberative assemblies often restrict attention to situations with only two options. Unfortunately, with a binary choice, the Diversity Trumps Ability Theorem is vacuous because, as I shall now show, it is not possible to simultaneously satisfy the assumptions imposed on the set of admissible search rules.

To see why, suppose that $X = \{\bar{x}, x^*\}$ with $V(\bar{x}) < V(x^*)$. There are four possible search rules, as shown in Table 2. Neither ϕ^b and ϕ^c satisfy Assumption 1.(1) because if the search starts at x^* , it does not terminate there. Because ϕ^b always chooses the other alternative wherever it starts, it also fails to satisfy Assumption 1.(2). Because the problem is difficult, by Assumption 2, ϕ^d cannot be in Φ because it always chooses x^* . Thus, the only possible choice for Φ is $\{\phi^a\}$. But with ϕ^a , the search always stops where it starts, so there is no diversity, violating Assumption 3. Hence, Φ is empty, thereby establishing Theorem 2.¹⁸

Theorem 2. *When the set of alternatives X is binary, it is not possible to satisfy all of the assumptions of the Diversity Trumps Ability Theorem because no nonempty set of search rules Φ can simultaneously satisfy Assumptions 1–3.*

In terms of democratic decision-making, Theorem 2 shows that the Diversity Trumps Ability Theorem is irrelevant when decisions are binary; it cannot be used to make the case that diversity does a better job of tracking the truth than expertise. In such situations, the Diversity Trumps Ability Theorem offers no comfort to those who want to use it to argue for the collective decision to be made by an inclusive set of individuals rather than by an epistocracy.

¹⁸Note that the emptiness of Φ has been established without considering Assumption 4.

7. Non-Binary Decisions

Although the assumptions of the Diversity Trumps Ability Theorem are inconsistent for binary decisions, they are not when there are more alternatives. However, even in non-binary choice situations, there are grounds for caution in drawing strong implications for the design of democratic institutions from this theorem. To use a distinction employed by Landmore (2014), these concerns can take one of two forms: internal or external. In the former case, the implications of this theorem are evaluated in terms of the internal coherence of the arguments used taking the truth-tracking objective as given. In the latter case, other objectives are also considered.¹⁹

Anderson (2006) employs the external perspective when discussing the shortcomings of using the Diversity Trumps Ability Theorem to legitimize an epistemic democracy. According to her, this theorem provides an exclusively instrumental view of democratic deliberation and participation, thereby failing to account for the intrinsic importance of universal inclusion. Moreover, it does not take account of the role that periodic elections and other feedback mechanisms play in enhancing the predictive performance of democratic decision-making over time.

The political application of the Diversity Trumps Ability Theorem can also be evaluated internally. One could, for example, determine to what extent each of the theorem's assumptions apply to democratic assemblies. However, the assumptions of the theorem are merely sufficient for its conclusion, and this conclusion may be fairly robust to changes in some of the assumptions used, so such an exercise may be of limited value. Nevertheless, there are some assumptions that are quite fundamental, so it is worth considering them.

The Diversity Trumps Ability Theorem only shows that the epistemic advantages conferred by diversity are realized in a sufficiently large population. In principle, the requisite size could be extremely large. As Landmore (2013) recognizes, the practical difficulties of deliberating in large groups can be substantial, so if this theorem is to have practical relevance in non-binary situations, the size of the deliberating group cannot be too large. Hong and Page (2004) report experimental evidence suggesting that the benefits from diversity materialize in fairly small groups, which is suggestive, but far from

¹⁹See Peter (2009) for an extensive discussion of the objections to basing an assessment of democratic institutions solely on their ability to track the truth well.

conclusive, that unrealistically large groups are not needed to apply this theorem.

As part of her argument for the inclusiveness of deliberative democratic institutions, Landmore (2013, p. 104) conjectures that the Diversity Trumps Ability Theorem can be generalized to a Numbers Trump Ability Theorem because larger groups are more likely to be cognitively diverse. But how well a small or medium-sized group tracks the truth depends not just on its being cognitively diverse, it also depends on the exact nature of the diversity. For example, if an individual with a novel search rule is added to a moderately-sized group, he might direct the group to an outcome somewhat further from the best one because there is not enough diversity for anybody else to recognize that a further improvement is possible. Moreover, enlarging the number of decision makers may have no impact on their epistemic competence if the added search rules are similar to ones that are already being used. This suggests that a more nuanced approach to capturing the benefits of diversity in reasonably sized groups is needed.²⁰

The Diversity Trumps Ability Theorem presupposes that any information that is shared is communicated truthfully. In particular, each individual truthfully reports any local optimum that he finds during his search process. However, while everybody shares the same goal of identifying the best alternative, as a result of their searches, individuals have different beliefs as to what this outcome is. As first recognized by Austen-Smith and Banks (1996) in the context of the Condorcet Jury Theorem, this difference in beliefs may provide incentives for individuals who share a common goal to misrepresent what they know because they think that this will further the group's interest in finding the best outcome.²¹ Indeed, experts on an issue may well regard the reported local optima of those who they feel are uninformed to be unreliable, and so misreport their own local optima in order to encourage convergence on an outcome that they think is closer to the truth. In her discussion of this problem as it applies to the Condorcet Jury Theorem, Lan-

²⁰For example, the approach used by Guerdjikova and Nehring (2014) to aggregate the opinions of a group of individuals into an overall judgment takes account of the dissimilarity of their opinions. In predictive contexts, Landmore and Page (2014) note that the correlation between the predictive models used matters, not just the number of models. The same can be said of search rules in the present context.

²¹List and Pettit (2011) provide a good introduction to this incentive issue. A brief overview of the literature dealing with this issue may be found in Bozbay, Dietrich, and Peters (2014).

demore (2013, p. 156) notes that strategic considerations play no significant role in large groups because in such settings any single individual has little influence on the outcome. But for realistic sizes of deliberative assemblies, this need not be the case. Nor need it be the case that a small group of like-minded individuals will refrain from jointly misrepresenting what they have learned so as to influence the outcome in a way that they think better serves the group's interests than telling the truth.

The idea that diversity can enhance the collective wisdom of a group is a powerful one. But, as I have argued, in binary choice situations, the formalization of this idea in the Diversity Trumps Ability Theorem provides no grounds for favouring diversity over expertise when the objective is to identify the best outcome according to some procedure-independent standard of goodness. I have also suggested that when there are more than two options to choose from, there are reasons to believe that the theorem's assumptions about the size of the group and the behaviour of its members may not be satisfied by the kinds of democratic assemblies that Landemore has in mind. As a consequence, care must be taken when applying the Diversity Trumps Ability Theorem to democratic institutions.

References

- Anderson, E., 2006. The epistemology of democracy. *Episteme* 3, 8–22.
- Austen-Smith, D., Banks, J. S., 1996. Information aggregation, rationality, and the Condorcet Jury Theorem. *American Political Science Review* 90, 34–45.
- Bozbay, İ., Dietrich, F., Peters, H., 2014. Judgment aggregation in search for the truth. *Games and Economic Behavior*, forthcoming.
- Cohen, J., 1986. An epistemic conception of democracy. *Ethics* 97, 26–38.
- Estlund, D. M., 2008. *Democratic Authority: A Philosophical Framework*. Princeton University Press, Princeton, NJ.
- Guerdjikova, A., Nehring, K., 2014. Weighing experts, weighing sources: The diversity value, unpublished manuscript, Théma, Université de Cergy-Pontoise.
- Gunn, P., 2014. Democracy and epistocracy. *Critical Review* 26, 59–79.
- Hong, L., 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, 16385–16389.
- Landemore, H., 2013. *Democratic Reason: Politics, Collective Intelligence,*

- and the Rule of the Many. Princeton University Press, Princeton, NJ.
- Landemore, H., 2014. Yes, we can (make it up on volume): Answers to critics. *Critical Review* 26, 184–237.
- Landemore, H., Page, S. E., 2014. Deliberation and disagreement: Problem solving, prediction, and positive dissensus. *Politics, Philosophy and Economics*, forthcoming.
- LiCalzi, M., Surucu, O., 2012. The power of diversity over large solution spaces. *Management Science* 58, 1408–1421.
- List, C., Pettit, P., 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press, Oxford.
- Page, S. E., 2007. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton University Press, Princeton, NJ.
- Peter, F., 2009. *Democratic Legitimacy*. Routledge, New York.
- Quirk, P., 2014. Making it up on volume: Are larger groups really smarter? *Critical Review* 26, 129–150.
- Surowiecki, J., 2004. *The Wisdom of Crowds: Why the Many are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies, and Nations*. Doubleday, New York.