Christopher Thompson

A general model of a group search procedure, applied to epistemic democracy

Article (Accepted version)
(Refereed)

Original citation:
Thompson, Christopher (2013) *A general model of a group search procedure, applied to epistemic democracy.* Synthese, 190 (7). pp. 1233-1252. ISSN 0039-7857

DOI: 10.1007/s11229-012-0108-9

A general model of a group search procedure, applied to epistemic democracy

Christopher Thompson

(Forthcoming in: Synthese (Special Issue on the Epistemology of Inclusiveness). 23 February 2012 draft.

Abstract

The standard epistemic justification for inclusiveness in political decision making is the Condorcet Jury Theorem, which states that the probability of a correct decision using majority rule increases in group size (given certain assumptions). Informally, majority rule acts as a mechanism to pool the information contained in the judgements of individual agents. I aim to extend the explanation of how groups of political agents track the truth. Before agents can pool the information, they first need to find truth-conducive information. Increasing group size is also important in the initial search for truth-conducive information.

1. Introduction

At least some political decisions are of the kind that can be correct or incorrect. If it is the case that groups of political agents are more reliable at making correct decisions than individual agents, then this can provide strong epistemic grounds for widening the democratic franchise. Majority voting and the Condorcet Jury Theorem play a prominent role in most current accounts of epistemic democracy. Informally, judgement aggregation procedures such as majority voting serve as mechanisms to pool the truth-conducive information contained in the judgements of individual agents. The literature on the Condorcet Jury Theorem is already quite developed, and in this paper I mostly rehearse the main points. But what has largely been overlooked in the current literature is how political agents obtain truth-conducive information in the first place.

A small but interesting philosophical literature on search procedures has developed quite recently. Various authors have developed models of group search which, inter alia, consider the search behaviour of agents and the group composition which is epistemically optimal. I draw out the two types of social epistemic mechanism that allow groups of agents to be more successful in searches than individual agents. I show why it is the case that increasing the size of the group can increase the probability that an object of search will be found.

The epistemic justification for widening the democratic franchise can be expanded into a two-staged process. Firstly, agents engage in a search procedure to extract truth-conducive information from the environment. Increasing the size of the group increases the amount of information obtained by members of the group. Secondly, aggregation procedures such as majority voting allow the group to pool the information obtained by individual agents. Increasing the size of the group increases the amount of information contained in the social choice.

2. Judgement aggregation procedures

Some political decisions are about matters of fact. For example, whether a nation possesses nuclear weapons or not, which form of power generation has the lowest costs and which presidential candidate has the policies that will create the most jobs are all questions which have correct and incorrect answers. Propositions which describe possible states of the world are termed *'alternatives'*. Possible alternatives might include, for example 'that the nation in question does possess nuclear weapons', 'that the nation in question does not possess nuclear weapons'; 'wind power is cheapest', 'coal power is cheapest', 'nuclear power is cheapest', 'gas power is cheapest'; 'the Republican presidential candidate will create the most jobs' and 'the Democratic presidential candidate will create the most jobs'. To help interpret the votes of agents we often have an *agenda* which contains a specific set of alternatives. The agenda is common knowledge for all relevant parties. An agenda might contain a complete logical partition of possible states of the world such as 'that the nation in question does possess nuclear weapons'/ 'that the nation in question does not possess nuclear weapons'. It is possible that the agenda only contains some of the possible alternatives, for example 'wind power is cheapest'/ 'coal power is cheapest'. If the agenda only contains some of the possible alternatives then there is a risk that the correct alternative is not included. In what follows we assume that the agenda is comprised of two alternatives, and that only one of these alternatives is correct (only one of the propositions accurately describes the true state of the world)[i].

A political decision requires decision makers which are termed *'agents'*. Each agent (or voter, or juror) can express their judgement as to what they think the actual state of the world is, as to what they think the correct alternative is. Agents express their judgement by casting votes for particular alternatives. An aggregation procedure allows a group to generate a collective judgement (or social choice) which depends on the judgements of individual group members. It can be construed as "…a function which assigns to each combination of individual judgements across the group members a corresponding set of collective judgements" (List, 2008, p.289).

Figure 1: aggregation procedures (List, 2008, p.298)

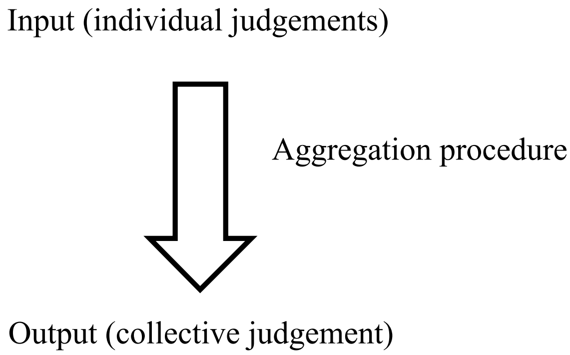

There are a variety of different aggregation procedures including (but not limited to) dictatorship, unanimity rule and majority rule. With dictatorship, the social choice is just the judgement of the single agent who is deemed the dictator. With unanimity rule, an alternative will be the social choice if and only if it receives the votes of all the agents. With majority rule, an alternative will be the social choice if and only if it receives strictly more than half of all the votes. These three aggregation procedures are particularly salient, and often feature in the literature on epistemic aspects of social choice theory[ii]. Each aggregation procedure has different virtues, but the focus of this paper is the epistemic virtue, the probability that an aggregation procedure will select the correct alternative (and avoid the wrong alternative[iii]) as the social choice.

Majority rule is the default democratic aggregation procedure. May's theorem shows that in a pairwise choice majority rule is the only aggregation procedure that satisfies the four important procedural virtues of universal domain, neutrality, anonymity and positive responsiveness[iv]. Majority rule also has important epistemic virtues. The epistemic potential of majority rule as a judgement aggregation procedure is supported by the classic Condorcet Jury Theorem[v] (CJT). The classic CJT applies to social choice problems in which simple majority voting is used to determine the social choice when there are two alternatives on an agenda, one of which is objectively correct.

The CJT has two assumptions:

- *Competence*: the probability that agents will vote for the correct alternative is homogeneous, greater than a half and less than certainty.
- *Independence*: the events of any two agents voting for the correct alternative are independent, conditional on the state of the world.

The classic CJT result comes in two parts:

- *Non-asymptotic CJT*: the probability that the group will select the correct alternative is monotonically increasing as the group size increases;
- *Asymptotic CJT*: in the limit as group size tends towards infinity, the probability of a correct majority verdict tends towards certainty.

The probability of a correct majority verdict (P) is as follows[vi]:

$$P = \sum_{h > \frac{n}{2}} \binom{n}{h} p^h (1-p)^{n-h}$$

where p is the homogeneous level of competence and n is the size of the group.

We can see the effect of increasing group size in the figure below. Here, the homogeneous competence level is $p = 0.6$ and agents are independent.

Figure 2: the probability of a correct majority verdict as group size increases.

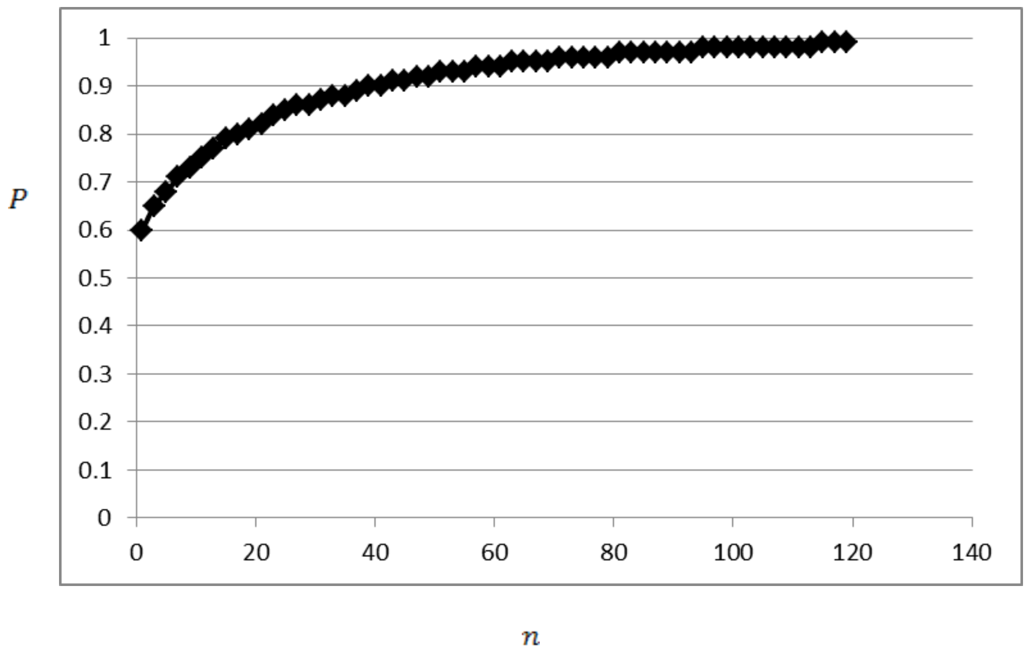

There have been various extensions of the classic CJT, two of which are of particular significance. Theorem V of Grofman et al. (1983) extends the CJT to heterogeneous competencies, where the average competence level is greater than a half and the distribution of competencies of the group members is symmetric. List and Goodin

(2001) extend the CJT from majority rule on a two-placed agenda, to plurality rule on a many-placed agenda.

It should be noted that the ability of the judgement aggregation procedure of majority rule to track the truth goes beyond the limited conditions of the CJT and its extensions. For example, there will be some cases where the distribution of competence levels in a group is *not* symmetric about the mean. Nevertheless the probability of a correct majority verdict may be greater than the competence level of any single member of the group and may in fact be close to certainty. Similarly, there may be cases where the independence assumption of the CJT is violated and yet the probability of a correct majority verdict is extremely high. However, if the two assumptions of the CJT are not met increasing group size is not always epistemically virtuous, we have no general rule for determining the epistemic impact of increasing group size and instead we need to rely on sample calculations. For example, a group with competencies $(p_1, p_2, p_3) = (0.6, 0.6, 0.7)$ has a probability of a correct majority verdict of $P_{Majority\,Rule} = 0.696$. If we are concerned with making the correct social choice then in this case it would have been better to make agent 3 the dictator, and ignore the votes of the other two agents. However a group with competencies $(p_1, p_2, p_3, p_4, p_5) = (0.6, 0.6, 0.7, 0.6, 0.7)$ has a probability of a correct majority verdict of $P_{Majority\,Rule} = 0.75$. By employing majority rule this group of five agents is more reliable than any individual group member.

How are we to account for the epistemic power of majority rule? Formal proofs of the asymptotic CJT have been published elsewhere[vii]. Here I will present two informal explanations for why majority rule can be successful at tracking the truth.

Firstly, majority rule has a tolerance for mistaken, incorrect votes. With the judgement aggregation procedures of dictatorship or unanimity rule, if a single agent makes the mistake of voting for the incorrect alternative, the correct alternative will not be the social choice. By contrast, with majority rule the mistaken votes of a minority of agents can be off-set by the correct votes of the majority. If the competence assumption of the CJT holds then as group size increases it becomes increasingly likely that there will be more votes for the correct alternative than for the incorrect alternative (in line with the law of large numbers). Suppose, for example, that the group was comprised of one agent with competence $p_1 = 0.6$. This agent has a 0.4 probability of making a mistake and voting for the incorrect alternative. However if we increase the group size to three agents with competencies $(p_1, p_2, p_3) = (0.6,0.6,0.6)$ and employ majority rule, then when one agent makes the mistake of voting for the wrong alternative, their error can be compensated for by the correct votes of the other two agents. Under majority rule, the correct alternative will still be the social choice even if any one of the three agents makes a mistake and votes for the incorrect alternative. The probability of correct majority decision, given the competencies in this group, is $P_{Majority\,Rule} = 0.648$.

The trade-off between correct and incorrect votes can only occur if there is some diversity in the voting behaviour of agents. The diversity in voting behaviour by agents is captured by the independence assumption of the CJT. If agents were entirely dependent, and voted identically, then there would be no epistemic advantage to increasing group size. For majority rule to track the truth as group size increases, it must be possible for an agent to vote correctly if another agent votes incorrectly.

The second informal explanation for the truth-tracking ability of majority rule is that judgement aggregation procedures, like majority rule, act as information pooling mechanisms. As Ladha, 1992, states: "…the majority-rule mechanism is simply a means to aggregate the experts' opinions and thereby their information… when the assumptions of the CJT are met, majority rule voting serves as a mechanism to assimilate decentralised information about the alternatives." (p.619)[viii]. Each agent has an incomplete view of the state of the world. By combining these incomplete views, the group is able to pull together a complete view of the true state of the world.

If majority rule is to pool the information contained in the judgements of individual agents, it must be the case that the judgements contain some truth-conducive information. The competence of agents, p, represents the probability of the event that agents vote for the one correct alternative on a two-placed agenda. If an agent had no information whatsoever as to which alternative on the agenda is correct, and is simply casting a vote at random, then he or she will have a competence of $p = 0.5$. If an agent is to have a competence level better-than-random, this agent needs to receive a truth-conducive piece of information which makes him or her more likely to vote for the correct alternative than the incorrect alternative. The competence of this agent, given the true state of the world and some truth-conducive information, will be $p > 0.5$.[ix]

The competence assumption of the classic CJT can be weakened – it does not have to be the case that each agent receives a piece of information such that the competence levels of all agents in the group are identical. It is sufficient for each agent to receive information of differing truth-conducive strengths such that the average competence level is greater than ½, and the competencies of group members are symmetrically distributed[x].

The independence assumption of the classic CJT requires that the vote of one agent (for the correct or incorrect alternative) makes it neither more nor less likely that another agent votes for the correct or incorrect alternative. If agents have the same piece of truth-conducive information in common, then the independence assumption will be violated. The fact that one agent votes correctly will make it *more* likely that another agent will votes correctly (since the common piece of information serves as a common causal factor on their votes). If agents have no information in common, then the vote of one agent should make it neither more nor less likely that another agent will vote correctly. The independence assumption of the CJT requires that as each new agent is added to the group, the agent brings with them entirely *new* pieces of truth-conducive information.

It may be possible to weaken the independence assumption of the classic CJT. Ladha (1992) and Estlund (1994) each consider the impact of violations of the independence assumption, of shared information and correlated votes. Ladha argues that the probability a majority verdict is correct is inversely related to the average of the coefficients of correlation. In other words, if some of the information which generates the competence levels of agents (and which is pooled by the aggregation procedure of majority rule) is held in common between agents, then the probability of a correct majority verdict will be lower. However, provided that agents have at least some private truth-conducive information (provided the votes of agents are not entirely dependent) there may be epistemic advantage to increasing the size of the group. Estlund argues that the presence of common influences does not easily rule independence in or out, and in fact deference to more competent opinion leaders can be epistemically virtuous.

In sum, if majority rule is to track the truth as group size increases, the agents added to the group need to bring with them new pieces of information. But agents do not have this information a priori. Current epistemic accounts of democratic decision making, which rely on majority rule and the CJT, only start part way through the process. The current accounts show how information contained in the judgements of agents is pooled into the social choice, but the current accounts are silent on how agents obtain this information. To complete the epistemic account of democratic decision making we need an account of how agents extract information from the environment in the first place.

Increasing group size is epistemically virtuous for the judgement aggregation procedure of majority rule, provided that agents are minimally competent (as per the CJT competence assumption) and provided there is some diversity in the voting behaviour of agents (as per the CJT independence assumption). In particular it must be possible for an agent to vote correctly when another agent votes incorrectly, so that the correct votes can off-set the minority of incorrect votes. Increasing group size may also be important in the search for information. As with the aggregation procedure of majority rule the success of a group at finding truth-conducive information may rely on differences in the behaviour of agents. With the judgement aggregation procedure of majority rule, it is differences in the voting behaviour of agents that is important. With search procedures, it may be differences in way agents search for information that makes increasing group size epistemically virtuous. If there are some differences in the search behaviour of agents, then the failure of one agent to find a piece of information can be compensated for by another agent.

3. Group search procedures

A search procedure can be construed as a function which assigns to each agent a corresponding set of objects. Individual agents have their own search procedure. A group's search procedure is comprised of the search procedures of the individual agents in the group and the success of a group at finding objects depends on the success of the individual search procedures. A group search procedure allocates subsets of objects from the total search space to individual members of the group.

Figure 3: search procedures.

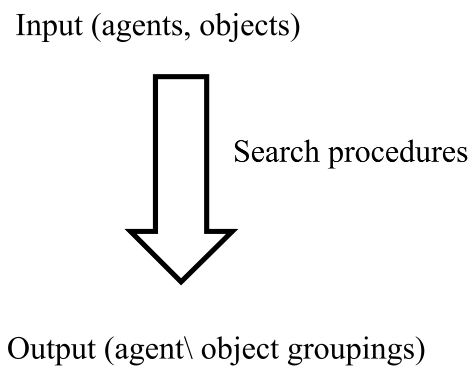

The idea of a search procedure should be familiar. Suppose you are writing a paper on Wittgenstein. You need to find the specific passage where Wittgenstein states that we perceive a proposition as a picture. To find the quote you need to engage in a search – to move from one possible location of the quote to the next, to see if that is where the quote is located. Suppose you have looked through the Tractatus and the

Philosophical Investigations, but just can't find the passage you are looking for. Your chances of finding the quote will be improved if you email your colleagues asking for help. Provided your colleagues have at least some chance of finding the quote (provided they are familiar with the Tractatus and Philosophical investigations) and provided there are some differences in the way they search for the quote, then as the number of people searching for the quote increases, the chances that at least one person will find the quote for you should increase.

There are two possible mechanisms that can account for the success of the group search procedure. Firstly, given time constraints, it may not be possible for you or any of your colleagues to search every possible location for the quote. You may only have enough time to search points 1 and 2 of the Tractatus. Your colleague may only have enough time to search points 2, 3 and 4 of the Tractatus. Provided that there is some diversity in the locations visited by agents, then as the number of agents increases, the total number of locations visited increases and the probability that someone finds the quote increases. In the limit, all possible locations for the quote will be visited by at least one agent and the quote is certain to be found. I term this type of group search procedure a 'spatial search procedure'.

The second possible mechanism that can account for the success of the group search procedure is that each agent might be fallible at recognising the quote when they visit the location of the quote. Suppose you read point 4.012[xi] but fail to take in its significance. A colleague also reads point 4.012 but fails to take in its significance. Provided that there is some diversity in the ability of the different agents to recognise the quote, then as the number of agents visiting the location of the quote increases, the chance that at least someone will recognise the quote and take in its significance increases. In the limit, someone is bound to recognise the quote eventually. I term this type of group search procedure a 'search recognition procedure'.

A small but interesting philosophical literature on search procedures has developed quite recently. Weisberg and Muldoon (2009) provide a spatial search model to investigate the division of cognitive labour in communities of scientists working in the same research field. Weisberg and Muldoon consider 'control' agents, who ignore the research of other agents, 'follower' agents, who move towards the research of other agents, and 'maverick' agents, who move away from the research of other agents. Using Netlogo computer simulations, Weisberg and Muldoon show that where the make-up of the group is homogeneous[xii], the proportion of the search space explored by the group increases as the size of the group increases. Importantly, they also show how it is epistemically desirable, from the groups' perspective, to have a mixture of 'maverick' and 'follower' agents in the search for successful approaches to particular scientific research topics. The maverick agents strike out on their own, away from the research of others, to find research areas of epistemic significance. Follower agents move towards the discoveries of other agents and help fully explore the areas of epistemic significance identified by maverick agents. Although the authors focus on the spatial aspect of search ('exploration' in their terminology), they also point to a search recognition aspect to the division of scientific cognitive labour[xiii].

Hong and Page (2004) also produce a spatial model of group search behaviour. Through proofs and computer simulations they show that 'diversity trumps ability', that a group with varied but sub-optimal search heuristics will outperform a group with optimal but similar search heuristics. The intuition behind their surprising result is roughly that the more varied a set of search heuristics, the more thoroughly a search space will be investigated and the more likely it is that the objects of interest (or 'solutions') will be identified.

List et al. (2008) produce a model of nest site selection by hives of honey bees. Their model combines aspects of a spatial search procedure, search recognition procedure and judgement aggregation procedure. List et al. stress the important interplay between interdependence and independence in the search behaviour of the bees. The bees need to be interdependent in their spatial search: a honey bee is more likely to visit a potential nest site if the location of that site is communicated to them by other bees. If the judgements of honey bees regarding the best nest sites are independent, then they will contribute new information to the group.

I will now articulate, in more precise terms, the two different types of search procedure: the spatial search procedure and the search recognition procedure. I present these as conjectures. Franz Dietrich helped me correctly formulate these conjectures.

Suppose we have a set of objects which are the subject of search. The objects of search might be the statements by Wittgenstein that compare propositions to pictures, or the objects might be truth-conducive pieces of information that indicate which presidential candidate will generate the most jobs. We have a finite set of locations where the object might occur. For example, the Wittgenstein quote could occur anywhere between points 1 and 7 of the Tractatus (or alternatively anywhere between page 1 and page 111 of the Tractatus). The information indicating which presidential candidate is best might be located in speeches given by candidates, or in press releases, or in articles written by commentators. The set of locations can be divided into jointly exhaustive (though perhaps not exclusive) subsets of locations. Each of these subsets of locations is visited by a different agent.

Each object of search occurs at a particular location[xiv]. The mapping from the set of all objects to the set of locations is - initially - unknown to agents in the group. If an agent moves to the location of an object the agent will have a certain probability of recognising that object.

We will consider the special case in which there is just one object which is the subject of search[xv]. If we can show that the probability of finding one particular object increases with group size, and it is the case that the ability to find objects is similar for all other objects, then it follows that as the size of the group increases the total number of objects found by the group increases. The location of the object and the subset of locations in which it occurs are initially unknown to the group. Each agent is assigned one subset of locations. Being assigned a subset of locations means that an agent visits each location in that subset as part of their search for the object. Once the agent moves to a location in their subset of locations, they attempt to identify the objects located there.

I will begin by presenting the search conjectures separately, before providing a combined conjecture.

<u>Spatial Search Conjecture</u>

Initially we will assume that each agent has an infallible ability to recognise an object if they visit the location of that object. For example, if an agent visits location 4.012 of the Tractatus, the agent is certain to recognise that this contains the quote the group is searching for. If the object occurs in the subset of locations visited by an agent it will be found by that agent.

We have two further assumptions, as follows:

Spatial Search Competence: For each agent, the unconditional probability that the object occurs in the subset of locations searched by an agent is uniformly bounded away from zero and is less than certainty.

Spatial Search Independence: The events of the object occurring in the subsets of locations searched by different agents are independent.

The probability that a member of a group of $1,2,\ldots,n$ agents finds an object of interest:

- *(non-limit claim)* increasing in group size n; and
- *(limit claim)* in the limit approaches certainty.

If S_i is the event that agent i finds the object and P_S^n is the probability that some member of a group of n agents finds the object, then:

$$P_S^n = \sum_{i=1}^{n} \Pr(S_i)$$

The Spatial Search Competence assumption of the Spatial Search Conjecture allows for agents in the search group to have heterogeneous probabilities of finding the object of search. The independence assumption in the spatial search conjecture reflects the need for there to be some differences or diversity in the way in which agents move around the search space looking for the object. If all agents visited exactly the same locations, then there would be no epistemic advantage to increasing the size of the group. The independence assumption reflects a neutral interpretation of the diversity in search behaviour. This assumption can be both strengthened and weakened. The strongest requirement for diversity in the spatial search behaviour of agents is as follows:

Spatial Search Diversity: For any two agents the events of the two agents visiting the location of the object are mutually exclusive. No agents have any locations in common and so it is impossible for two agents to find the object.

Of course the assumption that the sets of locations visited by agents are exclusive is quite demanding. If this assumption were to hold in practice, then it would require either a social planner to divide up the search space into non-overlapping subsets of locations; or it would require agents to communicate in the partitioning of the search space.

The Spatial Search Diversity assumption can be weakened. We can allow that there is some overlap in the locations visited by agents and therefore that the probabilities of different agents finding the objects are not independent. The minimum amount of private search we require from agents can be characterised as follows:

Spatial Search Diversity 2: Although the intersection in the set of locations searched by two agents may be non-empty, each agent has at least some locations that they search privately.

<u>Search Recognition Conjecture</u>

Here we assume that each agent in the group is certain to visit the location of the object of search, but agents have a less than perfect ability to recognise the object located there. For example, agents would be guaranteed to visit location 4.012 of the Tractatus, but agents are not guaranteed to recognise the significance of the quote there.

We have two further assumptions as follows:

Search Recognition Competence: The conditional probability that an agent recognises the object, given the object is in the set of locations searched by the agent, is uniformly bounded away from zero and less than certainty.

Search Recognition Independence: The events of different agents recognising the object, given that they visit the object's location, are independent.

The conditional probability that one of a group of agents recognises the object, given the object is in the set of common locations visited by the agents, is:

- (limit claim) increasing in group size n; and
- (non-limit claim) in the limit tends to certainty.

We write $R_i|S_i$ for the event that agent i recognises the object, given the agent visits the object's location. The conditional probability that a group of n agents recognises the object at a particular location, P_R^n, is given by:

$$P_R^n = \sum_{i=1}^{n} \Pr(R_i|S_i)$$

There may be some violations of Search Recognition Independence. An agent's ability to recognise objects could be caused by any number of factors. For example, an agent's ability to recognise Wittgenstein's quotes could be caused by the seminars or tutorials they attended which focussed on particular aspects of Wittgenstein's work. If two agents share some recognition ability generating factors (if, for example, they attended the same seminars) then their recognition abilities will not be independent. The probability of an agent recognising an object, given that their colleague has recognised the object, will be greater than the agent's unconditional probability of recognising the object. However independence in object recognition ability is secured by conditionalising on common factors as follows:

Search Recognition Independence 2: The events of different agents recognising the object are independent, conditional on the object being contained in the common set of locations and on factors held in common between agents.

Combined Search Conjecture

We can combine the two conjectures. If the Spatial Search Competence, Spatial Search Independence, Search Recognition Competence and Search Recognition Independence assumptions hold then the unconditional probability that one of a group of agents finds the object:

- (limit claim) increasing in group size n; and
- (non-limit claim) in the limit tends to certainty.

The probability that a group of n agents recognises the object at a particular location, P_F^n, is given by:

$$P_F^n = \sum_{i=1}^{n} \Pr(S_i) \times \Pr(R_i|S_i)$$

The two mechanisms driving the epistemic performance of the Combined Search Conjecture are, firstly, that different agents visit different locations (Spatial Search Conjecture); and, secondly, that different agents visit the same location but have differing abilities to recognise the object located there (Search Recognition Conjecture). These two mechanisms pull in different directions. If we encourage agents to disperse and visit different locations we decrease the probability that the objects at those locations will be recognised. If instead we encourage agents to visit the same locations we increase the probability the objects at that particular location will be recognised, but we decrease the probability of finding objects that occur at different locations.

There are similarities in the mechanisms that account for the epistemic success of the two search procedures and the judgement aggregation procedure of majority rule. A single agent may fail to find the object of search. However provided that agents have some capacity to search for the objects (as per the search competence assumptions) and provided there are some differences in the search behaviour of agents (as per the search independence and diversity assumptions), then the new agents added to the group can compensate for failure of other agents to find the object.

Simulations

The search procedures were reproduced in the computer program NetLogo 4.1[xvi]. The assumptions of the search conjectures were satisfied and the model was tested to see if it confirmed the conjectures. Following the NetLogo conventions, the search space (set of all possible locations[xvii]) is represented in a two dimensional x and y tortoidal grid. The grid is 37 locations wide and 37 locations tall meaning that there are a total of 1369 locations in the search space. Each agent is placed at a random starting location and then moves from location to location according to a search heuristic. For example an agent's search heuristic may require the agent to rotate a random number of degrees to the right, and then move forward one location. Each simulation lasts 100 agent moves.

Firstly, I present the simulation results for the Spatial Search Conjecture where both the Spatial Search Competence and Spatial Search Independence assumptions hold.

Secondly, I present sample calculations for the Search Recognition Conjecture. Finally I present simulation results for the Combined Search Conjecture.

In each simulation model the number of agents in the group was varied. The experimental result is the proportion of locations visited at the end of the 100 moves[xviii]. The object of interest could occur on any one of the 1369 locations. In the limit, if all the locations are visited, the object of interest is guaranteed to be found. Therefore as the proportion of locations visited by a group of agents increases the probability that the object will be found also increases.

For the sake of illustration, a screen shot of the first simulation model is seen in the figure below. Here ten agents were placed on the search space at random locations. The agents employed their search heuristic for 100 movements. The paths the agents took through the search space have been traced and the locations visited by agents change from grey to black.

Figure 4: a screen shot of the NetLogo simulation.

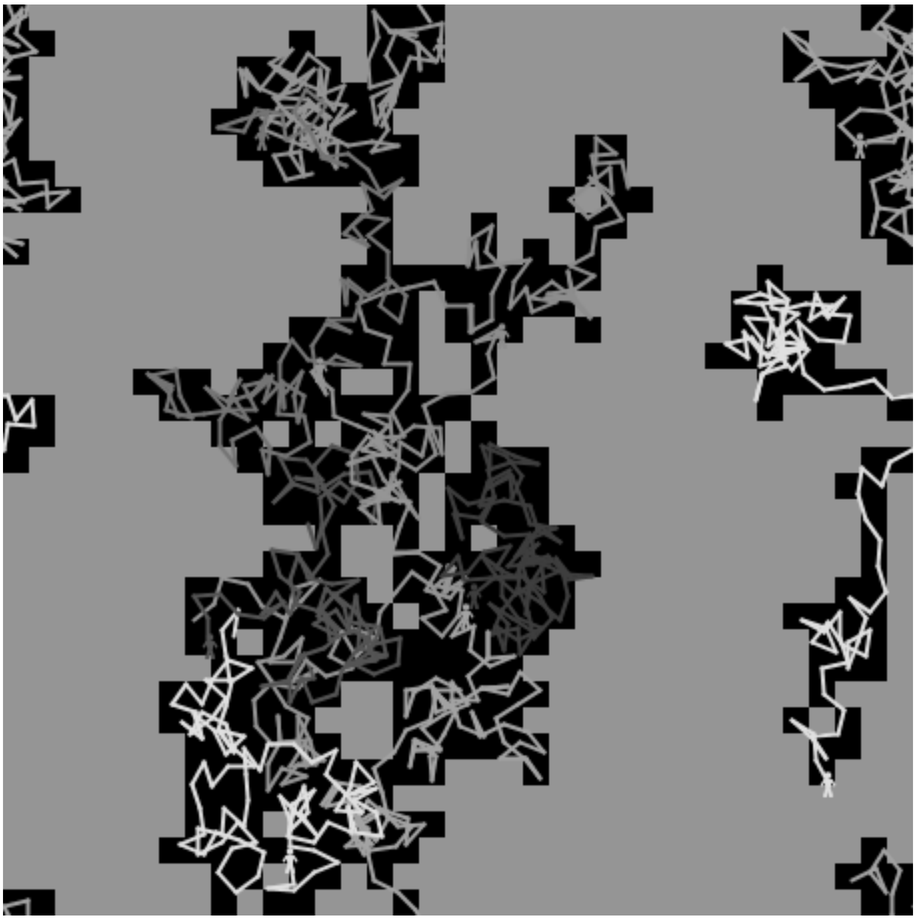

To assess the Spatial Search Conjecture in isolation, agents were placed at random start points on the search space. Here the Spatial Search Competence assumption holds since each agent in the group is placed on a location in the search space and any of these locations could contain the object of interest. Each agent in the group employs the same type of search heuristic whereby they rotate a random number of degrees to the right before moving forward one location. There is no restriction on agents exploring locations also visited by other agents, thus the Spatial Search Independence assumption holds. As can be seen in the figures below, the probability that an agent in the group will visit the location of the object (and by assumption find the object) is increasing and in the limit tends towards certainty. This simulation provides confirmation for the Spatial Search Conjecture.

Figure 5: a graph of the simulation results for the Spatial Search Conjecture

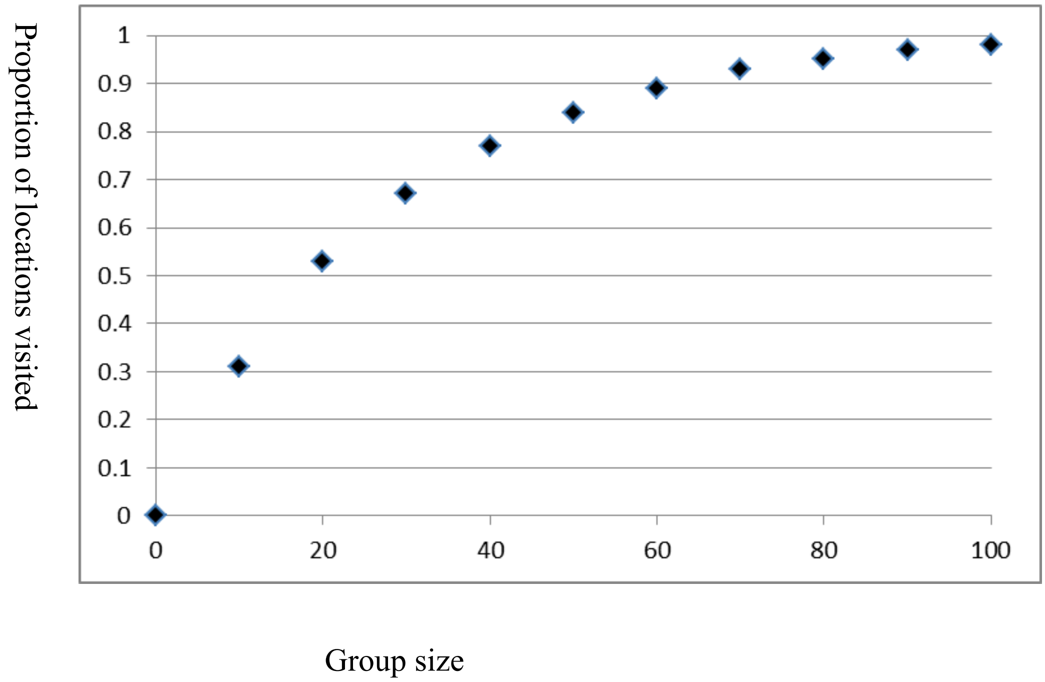

We now assess the Search Recognition Theorem in isolation via sample calculations. Here, to isolate the effect of additional agents on the probability of recognising the objects at particular locations, we assume that all the agents have reached the same location. Agents have a $\Pr(R_i|S_i) = 0.05$ probability of recognising the object if they move to that object's location – this is consistent with the Search Recognition Competence assumption. The events of agents recognising the object are independent, in line with the Search Recognition Independence assumption.

Figure 6: a graph of sample calculation results for the Search Recognition Conjecture.

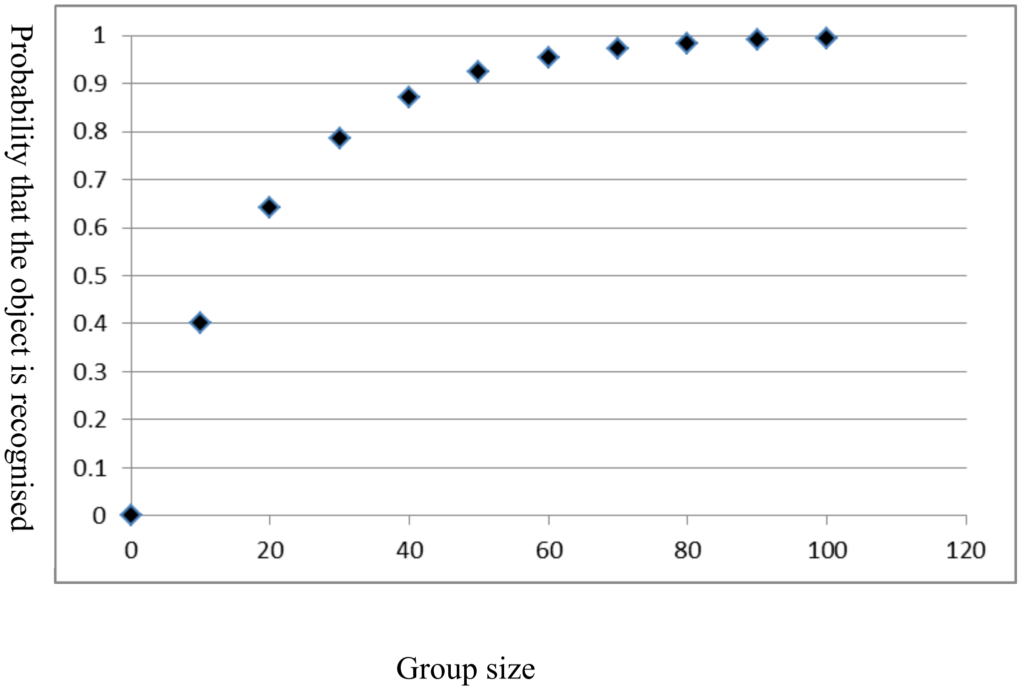

As can be seen, as group size increases, the probability that a member of the group recognises the object also increases. These sample calculations confirm the Search Recognition Conjecture. The main result to take away from these sample calculations is that even when recognition competence is low and even when the number of agents is small, the probability that at least one member of the group recognises the object will be high. Even when agents only have a 5% chance of recognising the object, if 50 agents visit that same location there is a better than 90% chance than at least one of the agents will recognise the object.

Finally I present the results of a simulation that models the Combined Search Conjecture. Here the start point of agents are determined randomly. Each agent in the group employs the same type of search heuristic whereby they rotate a random number of degrees to the right before moving forward one location. Spatial Search Competence holds under these circumstances. There is no restriction on agents exploring locations also visited by other agents. As such, Spatial Search Independence holds. We set agents level of recognition competence to 0.05 (and so Search Recognition Competence holds)[xix]. Finally, Search Recognition Independence holds since the event of one agent recognising the object makes it neither more nor less likely that another agent will recognise the object.

Figure 7: a graph of simulation results for the combined Search Conjecture (recognition competence $Pr(R_i|S_i) = 0.05$*).*

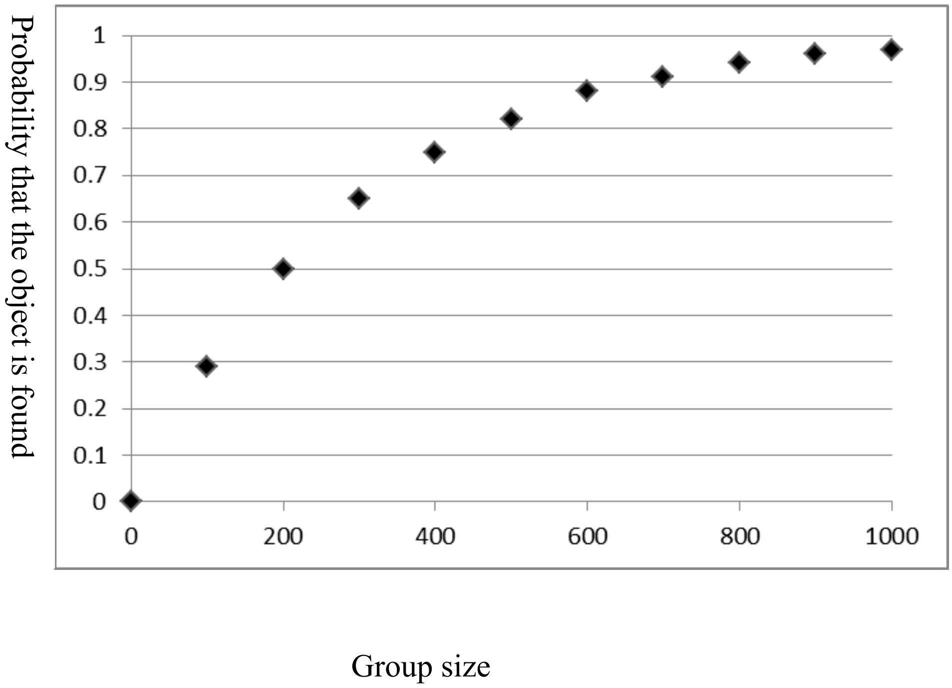

As can be seen in figure 7, as group size increases the probability that the object is found by at least one agent is increasing and tends towards certainty. This simulation result is in line with the Combined Search Conjecture.

4. Linking search and aggregation procedures

The discussion of search procedures in the previous section was general in nature. The two procedures of spatial search and search recognition could be applied to the search for any type of object, including tangible objects like a set of car keys you lost at home, and informational objects such as a Wittgenstein quote you cannot locate. The two search procedures can also be applied in a judgement aggregation framework. Recall that judgement aggregation procedures such as majority rule can only pool the information contained in the judgements of individual agents if in fact the judgements do contain some truth-conducive information. The competence of agents, p, represents the probability of the event that agents vote for the one correct alternative on a two-placed agenda. If an agent has no information whatsoever as to which alternative on the agenda is correct, and is simply casting a vote at random, then he or she will have a competence of $p = 0.5$. If an agent is to have a competence level better-than-random, this agent needs to receive a truth-conducive piece of information which makes him or her more likely to vote for the correct alternative than the incorrect alternative. The competence of this agent, given the true state of the world and some truth-conducive information, will be $p > 0.5$.

I will argue for a two-staged process for truth-tracking in groups of political agents. Firstly, there are search procedures by which agents extract truth-conducive information from the environment. Secondly, there are aggregation procedures (such as majority rule) which pool the information, dispersed across agents, into the social choice. Increasing group size is epistemically virtuous in each stage (given certain assumptions). During a search procedure, increasing group size increases the amount of information found by agents. During an aggregation procedure, increasing group size increases the amount of information pooled into the social choice. The success of both the search procedures and the aggregation procedure rely on agents having a minimum level of ‘competence’[xx] and there being some diversity in the behaviour of agents. For an aggregation procedure such as majority rule, there must be at least some chance of an agent voting correctly in the event of another agent voting incorrectly. For the search procedures there must be at least chance of an agent finding a piece of information in the event that another agent fails to find it. We can see a summary of the two-staged process in the figure below.

Figure 8: a two-staged process for truth-tracking in groups of political agents.

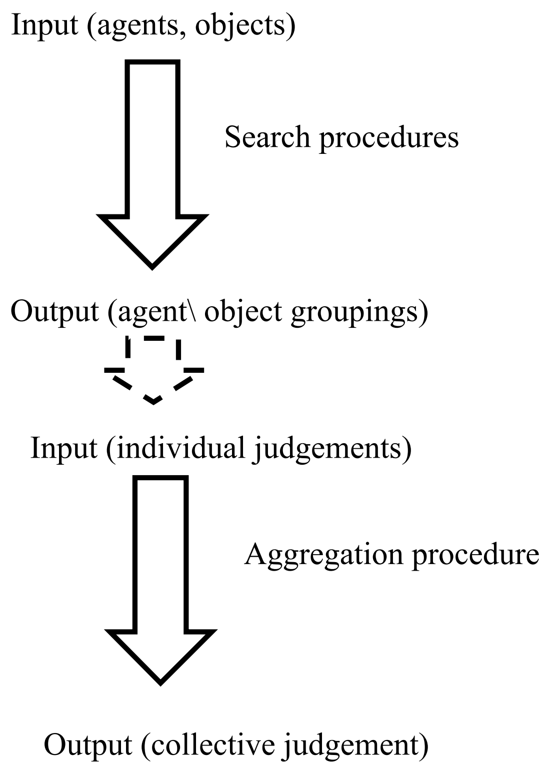

Consider the case in which, for the sake of argument, there is a consensus that the correct presidential choice is the candidate who will create the most jobs and where it is a matter of fact that one of the candidate will create more jobs than the other. The social choice (the candidate selected to be President) will be determined by majority rule. If the competence level of the agents (the voters) is to be better-than-random, if agents are more likely to vote for the correct than incorrect candidate, the agents need to receive some truth-conducive piece of information. If the votes of agents are to be independent, the agents need to receive different pieces of information. One such piece of truth-conducive information might be the fact that, as a state governor, one of the candidates presided over a period of job creation. A different piece of truth-conducive information might be the fact that candidates from a particular party have tended to have a stronger record of economic management than the other party. Each of these pieces of information may have a location. They may be located in manifesto documents, in speeches given by candidates, or in the assessments presented by pundits on news programmes. The more of this information captured by agents, and pooled into the social choice, the more likely the social choice is to be correct.

There may simply be too much truth-conducive information, indicative of the best presidential candidate for creating jobs, for a single agent to collect. Here we rely on a large number of voters to find the truth-conducive information in the first place, before sharing it with the group via their judgements. Firstly, it may be the case that several agents watch the same speech given by a candidate. However, the agents will not all pay attention to the same material in the speech. Agent's different capacities to recognise truth-conducive pieces of information mean that if one of the agents misses the part of the speech where the candidate talks of their success at creating jobs as a

Governor, another agent may pick up this statement. Secondly, it may be the case that the agents do not have the time to pay attention to all the potential sources of truth-conducive information. If one voter listens to the candidate speeches, while another reads manifesto documents, then a truth-conducive piece of information missed by one agent may be picked up by another agent.

5. Conclusion

I will conclude by pointing to two issues with the two-staged framework I propose, of search procedures followed by aggregation procedures, which mean that widening the democratic franchise is not always epistemically virtuous.

Firstly, the classic CJT only gives a conditional epistemic justification to widening the democratic franchise. The classic CJT states that *if* the competence and independence assumptions hold, then the probability of a correct majority verdict is increasing in group size and in the limit tends to certainty. If, on the other hand, the votes of agents are entirely dependent then increasing the size of the group will not increase the probability of a correct majority verdict. Even worse, if the probability of agents voting for the correct alternative on a two-placed agenda (their competence level) is *less* than ½ then as the size of the group increases the probability of a correct majority verdict *decreases* and tends to zero[xxi].

The competence and independence assumptions of the CJT will only be fulfilled if agents receive different pieces of truth-conducive information. For the agents to receive different pieces of truth-conducive information two conditions must be met.

Firstly, agents need to engage in a search for information (and the conditions of the search conjectures must be met). Secondly, the locations that agents search must contain some truth-conducive information. If there is no truth-conducive information to find then, post-search, agents will have no idea which of the two alternatives on the agenda is correct and their probability of voting for the correct alternative will remain at $p = ½$. If the locations searched by agents contain misleading pieces of information then by engaging in a search the agents become *less* likely to vote for the correct alternative. For example, it may be the case that a particular presidential election race is so polluted by misinformation that in searching for information before agents cast their votes the agents actually become *less* informed. There are three ways of addressing the problem of misleading information in the two-staged framework of search followed by aggregation. Firstly, the search recognition competence of agents could be expanded to include sensitivity to misleading information[xxii]. The revised search recognition competence of agents is then the probability that an agent recognises that a piece of information is relevant *and* that the information is truth-conducive (non-misleading), given that the agent moves to the location of the information. However, as Dietrich (2008) notes, a piece of information is misleading if it points to the incorrect alternative on the agenda. To know if a piece of information is misleading, we would need to know what the correct alternative is. But, by assumption, agents do not know what the correct alternative on the agenda is – that is why we are having a social choice exercise in the first place. So agents will simply be incapable of detecting whether a piece of information is misleading or not. A second approach to addressing the problem of misleading information is to limit the number of agents engaging in the search for information, so as to limit the amount of misleading information pooled into the social choice.

However, to be in a position to make the institutional decision to limit the group's search for information we would need to have at least some idea of which environments contain truth-conducive information and which environments are likely to be polluted with misleading information. Again, this may not always be an easy task. Finally, we could acknowledge that the possibility of misleading information really does pose a problem for epistemic justifications for widening the democratic franchise. As Dietrich and List (2004) prove[xxiii], the possibility of misleading information may mean that the maximum probability of a correct majority verdict is not certainty (as per the standard CJT) but rather a value much less than certainty.

The second of the two issues with the two-staged framework of search procedures followed by aggregation procedures is that the group's social choice is most likely to be correct if the judgement aggregation procedure employed by the group is sensitive to the post-search distribution of information across group members. If each agent has found some (different) truth-conducive information such that their competence levels are greater than ½, then majority rule may be the optimal voting rule. However, it is possible that the search space only contains a small pocket of truth-conducive information. Post-search it may be the case that almost all agents in the group have found no information and have a competence level of ½. A minority of agents may have found the truth-conducive information and have a competence level greater than ½. It may be epistemically advisable to increase the size of the group searching for information, so that we increase the chances that isolated pockets of information will be found by the group. However, it will be epistemically advisable to restrict the number of agents permitted to express their judgement to those agents with certain levels of competence. In some cases, if the competence level of agents is transparent, it may be epistemically optimal to employ a form of oligarchy such as expert dictatorship. Widening the democratic franchise is not always epistemically advisable.

Acknowledgements

I am grateful for the input from Richard Bradley, Christian List and Franz Dietrich. Franz Dietrich, in particular, provided important guidance in formulating the search conjectures. Ryan Muldoon provided helpful comments on an earlier version of this paper. I would also like to thank the anonymous referee for their suggestions. Versions of this paper were presented to the LSE Choice Group and the 3rd Copenhagen Conference in Epistemology: Epistemic Inclusiveness and Trust.

References

Bradley, R. and Thompson, C. (forthcoming, 2012) 'A (Mainly Epistemic) Case for Multiple-Vote Majority Rule', *Episteme.*

Dietrich, F. and List, C. (2004) 'A Model of Jury Decisions Where all Jurors Have the Same Evidence', *Synthese (special section: knowledge, rationality and action)*, 142(2): 175-202.

Dietrich, F. (2008) 'The Premises of Condorcet's Jury Theorem are not Simultaneously Justified', *Episteme*, 5(1): 56-73.

Dietrich, F. and Spiekermann, K. (unpublished 2010a), 'Epistemic Democracy with Defensible Premises'.

Dietrich, F. and Spiekermann, K. (unpublished 2010b), 'Independent Opinions?'.

Estlund, D.M. (1994) 'Opinion Leaders, Independence and Condorcet's Jury Theorem', *Theory and Decision*, 35(2): 131-162.

Estlund, D.M. (2008) *Democratic Authority: A Philosophical Framework*, Princeton University Press, Princeton, N.J.

Grofman, B., Owen, G. and Feld, S.L. (1983) 'Thirteen Theorems in Search of the Truth', *Theory and Decision*, 15(3): 261-278.

Hong, L and Page, S.E. (1994) 'Groups of Diverse Problem Solvers can Outperform Groups of High-Ability Problem Solvers', *Proceedings of the National Academy of Sciences of the United States of America*, 101(46): 16385-16389.

Ladha, K. (1992), 'Condorcet's Jury Theorem, Free Speech and Correlated Votes', *American Journal of Political Sciences*, 36: 617-634.

List, C. and Goodin, R.E. (2001) ‘Epistemic Democracy: Generalizing the Condorcet Jury Theorem’, *Journal of Political Philosophy*, 9(3): 277-306.

List, C. (2008) ‘Distributed cognition: a perspective from social choice theory’ in Albert, M. and Schmidtchen, D. and Voigt, S., (eds.) *Scientific competition: theory and policy*. Mohr Siebeck, Tübingen, pp. 285-308.

List, C., Elsholtz, C. and Seeley, T.D. (2008) ‘Independence and Interdependence in Collective Decision Making: An Agent-Based Model of Nest-Site Choice by Honeybee Swarms’, *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1518): 755-762.

May, K. O. (1952) ‘A set of independent necessary and sufficient conditions for simple majority decisions’, *Econometrica*, 20(4): 680–684

Owen, G., Grofman, B. and Feld, S.L. (1989) ‘Proving a Distribution-Free Generalization of the Condorcet Jury Theorem’, *Mathematical Social Sciences*, 17: 1-16.

Weisberg, M. and Muldoon, R. (2009) ‘Epistemic Landscapes and the Division of Cognitive Labour’, *Philosophy of Science* , 76(2): 225-252.

Wilensky, U. (1999) NetLogo. http://ccl.northwestern.edu/netlogo/. Centre for Connected Learning and Computer-Based Modelling, Northwestern University. Evanston, IL.

Wittgenstein, L. (1922) *Tractatus Logico-Philosophicus*, Kegan Paul, Trench, Trubner and co. Ltd.

---

[i] This is equivalent to there being one correct alternative on the agenda, with the other alternative on the agenda being a disjunction of anything not the correct alternative.

[ii] See, for example, List (2008) and Bradley and Thompson (2012).

[iii] I will also assume that the positive and negative reliabilities of agents are identical, that agents are just as able to correctly identify that a proposition is false as they are to correctly identify that a proposition is true.

[iv] May (1952). Universal domain means that any possible combination of votes are acceptable as inputs; neutrality means that the two alternatives on the agenda are treated equally, they each require strictly more than half of the votes to be the winner; anonymity treats all agents equally, each vote carries equal weight; positive responsiveness means that if the winning alternative x receives exactly one more vote than alternative not-x, and one agent were to change his or her vote from x to not-x, then alternative not-x would now be the winner.

[v] See, for example, Grofman, Owen and Feld (1983) Theorem I.

[vi] Owen, Grofman, and Feld (1989).

[vii] See, for example, Lahda (1992) and Dietrich (2008). I am unaware of any published proof of the full monotonic non-asymptotic CJT, but Dietrich and Spiekermann (2010a) includes such a proof. The CJT result relies on the law of large numbers, and is often explained by analogy with a sequence of coin tosses. See, for example, List and Goodin (2001) and Estlund (2008).

[viii] Similarly, Young (1995) states: "Condorcet showed that, if the voters make their choices independently, then the laws of probability imply that the choice with the most votes is the one most likely to be correct. In other words, majority rule is a statistically optimal method for pooling individual judgments about a question of fact." (p.52-53)

[ix] It is also possible that an agent receives a misleading piece of information which will make the agent less likely to vote for the correct alternative (more likely to vote for the incorrect alternative). Here, the

competence of the agent in question will be $p < 0.5$. For a more detailed taxonomy of the causal influences on an agent's vote (including informational causes) see Dietrich (2008).

[x] Grofman, Owen and Feld (1983) Theorem V.

[xi] "It is obvious that we perceive a proposition of the form *aRb* as a picture. Here the sign is obviously a likeness of the signified." Wittgenstein (1922).

[xii] Groups comprised solely of control agents, of follow agents or of maverick agents.

[xiii] Weisberg and Muldoon assume that if an agent discovers a significant scientific approach, the agent will be successful at fully exploiting its potential (making the important discoveries, getting the papers published). However, if the agents are not guaranteed to fully exploit the potential of a particular scientific approach then there may be group epistemic gains to be had from many agents following the same approach.

[xiv] For example, the Wittgenstein quote occurs at point 4.012 of the Tractatus.

[xv] For example, we are looking for the one quote where Wittgenstein states that we perceive propositions as pictures.

[xvi] Wilensky, U. (1999). The code for the simulations is based on the tutorial models provided by NetLogo with minor modifications. Code for the simulations is available on request. Note that Weisberg and Muldoon (2009) also use NetLogo in their simulations.

[xvii] Or 'patches' in NetLogo terminology.

[xviii] The experiment for each group size was run ten times, and the results reported are the average proportion of the locations visited.

[xix] In the simulation code, a location (or 'patch') has a 5% chance of turning from grey to black when visited by an agent, indicating that there is a 5% chance that any agent visiting the location will recognise the object there.

[xx] In a judgement aggregation framework, 'competence' is the probability an agent will vote for the correct alternative. In a spatial search procedure, 'competence' is the probability an agent moves to the location containing the object of search. In a search recognition procedure, 'competence' is the probability an agent recognises the object, given the agent moves to the location of the object.

[xxi] Grofman, Owen and Feld (1983) Theorem I.

[xxii] I am grateful to the anonymous referee for making this particular suggestion.

[xxiii] See also Dietrich and Spiekermann (unpublished a, b),