The London School of Economics and Political Science

Beyond Epistemic Democracy: The identification and pooling of information by groups of political agents.

Christopher Jeremy Thompson

A thesis submitted to the Department of Philosophy, Logic and Scientific Method of the London School of Economics for the degree of Doctor of Philosophy, London, September, 2011.

Declaration

I certify that the thesis I presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

Abstract

This thesis addresses the mechanisms by which groups of agents can track the truth, particularly in political situations.

I argue that the mechanisms which allow groups of agents to track the truth operate in two stages: firstly, there are search procedures; and secondly, there are aggregation procedures. Search procedures and aggregation procedures work in concert. The search procedures allow agents to extract information from the environment. At the conclusion of a search procedure the information will be dispersed among different agents in the group. Aggregation procedures, such as majority rule, expert dictatorship and negative reliability unanimity rule, then pool these pieces of information into a social choice.

The institutional features of both search procedures and aggregation procedures account for the ability of groups to track the truth and amount to social epistemic mechanisms. Large numbers of agents are crucial for the epistemic capacities of both search procedures and aggregation procedures.

This thesis makes two main contributions to the literature on social epistemology and epistemic democracy. Firstly, most current accounts focus on the Condorcet Jury Theorem and its extensions as the relevant epistemic mechanism that can operate in groups of political agents. The introduction of search procedures to epistemic democracy is (mostly) new. Secondly, the thesis introduces a two-stage framework to the process of group truth-tracking. In

addition to showing how the two procedures of search and aggregation can operate in concert, the framework highlights the complexity of social choice situations. Careful consideration of different types of social choice situation shows that different aggregation procedures will be optimal truth-trackers in different situations. Importantly, there will be some situations in which aggregation procedures other than majority rule will be best at tracking the truth.

Acknowledgements

Firstly, I would like to thank my supervisors Richard Bradley, Christian List and Franz Dietrich.

My development as a philosopher has also benefited from participation in the LSE Choice Group and interaction with various members of the group. Ryan Muldoon provided helpful feedback on the material contained in chapters 4 and 5 of the thesis. I also benefited from comments received on material presented at the 1st and 3rd Copenhagen Conferences in Epistemology, and the LSE-Groningen II exchange.

Above all, I would like to thank Jo. And Todd, Reg, Helen, Al and Nelly.

Table of contents

Chapter 1: Introduction. (p.11)

Chapter 2: Aggregation procedures. (p.41)

Chapter 3: The Condorcet Jury Theorem – agenda size and competence. (p.95)

Chapter 4: The generation of the inputs to aggregation procedures. (p.141)

Chapter 5: Group search procedures. (p.181)

Chapter 6: The link between search procedures and aggregation procedures. (p.229)

Chapter 7: The limits of the informational environment. (p.257)

Chapter 8: Conclusion. (p.311)

Bibliography. (p.319)

List of figures

Figure 1.1: a summary of the two-staged process of group truth-tracking. (p.24)

Figure 2.1: aggregation procedures. (p.43)

Figure 2.2: possible group judgements. (p.45)

Figure 2.3: possible individual judgements, given identical positive and negative reliabilities. (p.49)

Figure 2.4: the possible votes of a single agent. (p.70)

Figure 2.5: the majority winner, given a single agent. (p.70)

Figure 2.6: the probability of a majority winner, given a single agent. (p.71)

Figure 2.7: the possible combinations of votes, given three agents (p.71)

Figure 2.8: the majority winner, given three agents. (p.72)

Figure 2.9: the probability of a majority winner, given three agents. (p.73)

Figure 2.10: the possible combinations of votes, given five agents. (p.76)

Figure 2.11: the probability of a correct majority verdict, as group size increases. (p.80)

Figure 2.12: the probability of a correct majority winner, given dependent voters. (p.86)

Figure 2.13: the probability of $2,3's$ *votes.* (p.88)

Figure 2.14: the probability of a correct majority winner, given dependencies between agents. (p.88)

Figure 2.15: the probability of $2,3's$ *votes.* (p.89)

Figure 2.16: the probability of a correct majority winner, given some dependence between voters. (p.90)

Figure 2.17: a summary of the judgement-generating factors required for the aggregation procedures to track the truth. (p.93)

Figure 3.1: a quality ranking of alternatives. (p.97)

Figure 3.2: a mixed approach for multiple alternatives. (p.117)

Figure 3.3: a summary of the approaches for coping with multiple alternatives. (p.118)

Figure 4.1: an example of a complete causal network. (p.159)

Figure 4.2: an example of a causal network, from the perspective of an observer. (p.160)

Figure 4.3: a summary of the taxonomy of causal factors influencing an agent's vote. (p.161)

Figure 4.4: examples of causal networks where agents share causal factors. (p.168)

Figure 4.5: a revised taxonomy of causal factors. (p.171)

Figure 4.6: a causal network where each agent has private background and evidential factors. (p.176)

Figure 4.7: a causal network with common evidence, but private background factors. (p.176)

Figure 4.8: a causal network with common background factors, but private evidence. (p.177)

Figure 4.9: a causal network where a single agent receives all the information. (p.178)

Figure 4.10: a causal network with both private and common evidence and background factors. (p.179)

Figure 5.1: search procedures. (p.184)

Figure 5.2: a taxonomy of search mechanisms. (p.195)

Figure 5.3: a screen shot of the first simulation, before agents move. (p.207)

Figure 5.4: a screen shot of the first simulation, after 100 agent moves. (p.208)

Figure 5.5: simulation 1 results. (p.209)

Figure 5.6: graph of simulation 1 results. (p.209)

Figure 5.7: simulation 2 results. (p.210)

Figure 5.8: graph of simulation 2 results. (p.211)

Figure 5.9: a screen shot of simulation 2, after 100 agent moves. (p.212)

Figure 5.10: the probability an object will be recognised, as recognition competence varies. (p.214)

Figure 5.11: the probability an object will be recognised, as group size varies. (p.214)

Figure 5.12: simulation 3 results, recognition competence $Pr(R_i|S_i) = 0.5$. (p.215)

Figure 5.13: a graph of simulation 3 results, recognition competence $Pr(R_i|S_i) = 0.5$. (p.216)

Figure 5.14: simulation 3 results, recognition competence $Pr(R_i|S_i) = 0.05$ (p.217)

Figure 5.15: a graph of simulation 3 results, recognition competence $Pr(R_i|S_i) = 0.05$. (p.217)

Figure 6.1: the complete, two-staged, model of group truth-tracking. (p.233)

Figure 6.2: a taxonomy of the post-search spread of information across agents in a group. (p.252)

Figure 7.1: the probability of a correct social choice, given a premise-based or a conclusion-based procedure. (p.284)

Figure 7.2: a causal network with three private packages of information. (p.294)

Figure 7.3: a causal network with one common package of information. (p.295)

Figure 7.4: a causal network with three common packages of information. (p.297)

Figure 7.5: the probability of a correct majority verdict. (p.306)

Chapter 1: Introduction.

This thesis addresses the mechanisms by which groups of agents can track the truth, particularly in political situations.

I argue that the mechanisms which allow groups of agents to track the truth operate in two stages: firstly, there are search procedures; and secondly, there are aggregation procedures. Search procedures and aggregation procedures work in concert. The search procedures allow agents to extract information from the environment. At the conclusion of a search procedure the information will be dispersed among different agents in the group. Aggregation procedures, such as majority voting, then pool these pieces of information into a social choice.

The institutional features of both search procedures and aggregation procedures account for the ability of groups to track the truth and amount to social epistemic mechanisms. I identify two types of search procedure and three types of aggregation procedure whose respective institutional features are social epistemic mechanisms. Large numbers of agents are crucial for the epistemic capacities of each of these mechanisms. Interestingly, large numbers can be used in very different ways. We might task different agents in a group with performing the same task so that if some agents fail to perform the task other agents will be successful in performing the task. Instead we might task different agents in a group with performing different tasks so that the total number of tasks competed by the group will be large.

This thesis makes two main contributions to the literature on social epistemology and epistemic democracy. Firstly, most current accounts focus on the Condorcet Jury Theorem and its extensions as the relevant epistemic mechanism that can operate in groups of political agents. The introduction of search procedures to epistemic democracy is (mostly) new. Secondly, the thesis introduces a two-stage framework to the process of group truth-tracking. In addition to showing how the two procedures of search and aggregation can operate in concert, the framework also highlights the complexity of social choice situations. Careful consideration of different types of social choice situation shows that different aggregation procedures will be optimal truth-trackers in different situations. Importantly, there will be some situations in which aggregation procedures other than majority voting will be best at tracking the truth.

Background and limits of scope

I do not intend to give a comprehensive stand-alone literature survey for this thesis. Literature will be cited throughout the thesis whenever relevant. Here I will cite a few key texts to help place the thesis in the context of existing literature.

The thesis fits within the literature on social epistemology and epistemic democracy. Goldman (2010) provides a useful taxonomy of social epistemology. He notes that a variety of work streams go under the heading of ‘social epistemology’ and proposes a tripartite division of field. Firstly, social epistemology can focus on individual doxastic agents (IDAs) with social evidence. This aspect of social epistemology is the most continuous with traditional individualistic epistemology. Here the possessor of doxastic attitudes is still an individual agent but the sources of evidence for these attitudes are social in nature, such as the testimony of other agents. The second variety of social epistemology focuses on collective doxastic agents (CDAs). This departs from mainstream individualistic epistemology in that the possessor of doxastic attitudes is a group. The final type of social epistemology is systems-oriented social epistemology (SYSOR). As Goldman says, “An epistemic system is a social system that houses a variety of procedures, institutions, and patterns of interpersonal influence that affect the epistemic outcomes of its members.” (p.2). “In each case [for each social system] social epistemology would examine the systems in question to see whether its mode of operation is genuinely conducive to the specified epistemic ends. It would also identify alternative organizational structures that might be epistemically superior to the existing systems.” (p.8).

The focus of this thesis is very much on this third interpretation of social epistemology (SYSOR). The goal is not so much to assess whether the social institutions are successful at achieving their epistemic ends. Rather the goal is to identify the conditions under which social institutions *are* successful at achieving their epistemic ends and then account for *why* they are epistemically successful. Once we have this account of why the social institutions are epistemically successful we will be in a position to make normative claims, both concerning how the social institutions in question can be improved and how the social epistemic mechanisms identified can be applied to other settings.

There are many domains in which a social aspect of epistemology might be important. Goldman (1999) points to, inter alia, science, law, democracy and education as being significant. Providing an in-depth analysis of all the social epistemic mechanisms operating in each of these domains is beyond the scope of a single thesis. Instead this thesis focuses on the political domain. The choice of the term 'political' rather than 'democratic' is deliberate, as I do not want to exclude from consideration non-democratic political decision-making systems that may succeed at truth-tracking.

Almost all current literature on the topic of social epistemology as applied to political settings falls under the heading of 'epistemic democracy'. The term largely comes from Cohen (1986), though as Cohen points out, the idea that political decision making is at least in part about making *correct* decisions, and that different forms of government may be better or worse at making decisions has a distinguished history[1]. Plato's parable of the ship[2] suggests that the ship's owner (the citizenry) is bigger and stronger than anyone else on board but is deaf and short-sighted and has no knowledge of naval matters. The sailors (politicians) do not have the nautical skills to command the ship (the state), but

[1] Cohen (1986) cites Rousseau and Bentham. Wolff (1996) cites Plato, Rousseau and Mill.

[2] Plato (1998) s488a-189a.

compete with each other for the owner’s approval to take the rudder. If this doesn’t work, the sailors will subdue the owner, take over the ship and embark on a drunken voyage. Clearly for Plato there is a correct course that the ship of state should follow, and democratic forms of decision making are not suited for navigating this course. For Plato, statesmanship is a craft best carried out by philosopher kings who have the appropriate training.

Mill (1861) argues that the best form of government is one that has the greatest amount of beneficial consequences. ‘A completely popular government’ is the only form of government fitting this description since, inter alia, “… the general prosperity attains a greater height, and is more widely diffused, in proportion to the amount and variety of the personal energies enlisted in promoting it” (Chapter 3). Again we have the idea that political decisions can be correct and that a form of decision making that utilises the talents of the population is most likely to make these correct decisions. Mill seems to advocate a form of weighted majority rule (as discussed in the next chapter on this thesis): “When two persons who have a joint interest in any business differ in opinion, does justice require that both opinions should be held of exactly equal value?... One of the two, as the wiser or better man, has a claim to superior weight” (Chapter 8).

For Rousseau (1762) correct political decisions are those in line with the general will. The “…most natural arrangement is for the wisest to govern the multitude” (Book 3, Chapter 5), and the form of government that encapsulates this is an elective aristocracy.

Cohen (1986) produces an epistemic interpretation of voting, which is worth quoting in full:

> "An epistemic interpretation of voting has three main elements: (1) an *independent standard* of correct decisions — that is, an account of justice or of the common good that is independent of current consensus and the outcome of votes; (2) a *cognitive* account of voting — that is, the view that voting expresses beliefs about what the correct policies are according to the independent standard, not personal preferences for policies; and (3) an account of *decision making* as a process of the adjustment of beliefs, adjustments that are undertaken in part in light of the evidence about the correct answer that is provided by the beliefs of others." (p.34)

An epistemic populist, on Cohen's interpretation, argues that majority verdicts provide sound evidence about the common good (the independent standard of correctness). And the Condorcet Jury Theorem is frequently used to justify this claim.

The literature on epistemic defenses of democracy has been extended by several authors, including Estlund et al. (1989), Estlund (1997), Copp (1999), List and Goodin (2001), Anderson (2006) and Peter (2008). Majority voting and the Condorcet Jury Theorem play a prominent role in most of these accounts. Again, I do not intend to fully survey the path that research into epistemic democracy has taken over the last few years. Authors will be cited in the body of the thesis whenever relevant. Instead, I want to take the Cohen passage as a useful point of reference for characterising the scope of this thesis.

The core concern of this thesis is the mechanisms that operate in groups of political agents by which those agents can track the truth. For groups of agents to track the truth there must be an independent standard of correct decisions (as per Cohen's (1)). However, unlike Cohen, I do not limit the independent standard of correctness to the common good. There can be some standards of correctness that are independent of agent's preferences or judgements. For example, it may be false that a particular nation possesses nuclear weapons. Whether a group of agents believe that the nation possesses nuclear weapons, or whether a group of agents prefer that the nation possesses nuclear weapons, has no bearing on the fact that the nation does not possess nuclear weapons. There can be some independent standards of correctness that are the common good but where the judgements of agents only provide an imperfect indication of the common good. For example, it might be in the interests of everyone if the speed limit in urban areas were to be lowered to 40kph[3]. However, there may be considerable disagreement about this proposed policy and the votes of agents may be an unreliable guide to the common good. In line with convention I term the content of the social choice an 'alternative'. By assumption this thesis is only concerned with social choice problems in which there is one alternative that is objectively correct (with all other alternatives being incorrect). I generally focus on dichotomous choice problems where agents face an agenda

[3] For the sake of argument there would be fewer accidents, and traffic would flow more smoothly.

of two alternatives, one of which is correct, with the other alternative being incorrect. What counts as an alternative will vary according to the social choice problem. For example, if the social choice involves electing a political representative then the correct alternative might be the person uniquely qualified to be President. If the social choice problem involves policy choices over carbon-neutral power generation then the correct alternative might be nuclear power.

Given the reference point of a correct alternative I define three standards of group epistemic performance that will be of interest. Firstly, 'baseline epistemic performance' requires that a group of agents is better than random at selecting the correct alternative (or avoiding the incorrect alternative). Secondly, 'relative epistemic performance' requires that a group of agents is more likely than a single agent to identify the correct alternative (or avoid the incorrect alternative). Finally, 'absolute epistemic performance' requires that a group of agents is likely to select the correct alternative (or avoid the incorrect alternative). These three standards of epistemic performance are all important if we are to make the case that the institutional arrangements of a group mean the group is successful at tracking the truth. Furthermore if we are to make normative claims, on epistemic grounds, that a social institution such as majority rule should be implemented for political decision making then we should be able to show that a group using majority rule is more reliable than random at making the correct choice, more reliable than a single agent would be at making the correct social choice and likely simpliciter to make the correct social choice. ‘Tracking the truth’ is shorthand for meeting these three standards of epistemic performance.

Also important for this thesis is the cognitive account of voting (Cohen’s clause (2)). I focus on the beliefs expressed by agents and not their preferences. There is a large and interesting literature on the aggregation of preferences, but this is separate from epistemic issues. As Wolff (1994) has shown in his paper on the mixed motivation problem, if some agents vote according to their preferences while others vote according to their beliefs, then it is possible to have a social choice that is neither preferred by a majority nor believed true by a majority. There are two additional points that need to be made here. Firstly, I am primarily interested in the competence of agents, which is measured as the probability that they will vote for the correct alternative. To the extent that beliefs are discussed, they are characterized by the binary ‘believe an alternative is correct’ or ‘do not believe and alternative is correct’, as reflected in agents’ votes for or against an alternative. Fine grained degrees of belief and their translation into voting behaviour are not discussed. Secondly, the thesis is not only concerned with voting behaviour of agents. As will be argued, there are more epistemic mechanisms operating in groups of political agents than merely the aggregation of judgments. Searching for information in the first place and discussion between agents are also important. So the cognitive account of voting needs to be expanded into an account according to which agents are interested in tracking the truth and only express their opinions or judgments (not their preferences). We can further specify that agents do not engage in strategic behaviour. For example, they will not knowingly express incorrect views if they believe that by doing so the group as a whole is more likely to make a correct decision.

Finally, the third component of Cohen's conception – an account of decision making – will be deemed necessary but not sufficient for an account of political epistemic mechanisms. If the group is to track the truth it will need a final judgment as to the correct alternative. However, we need more than this. As will be argued, an important part of an epistemic conception of democracy is an account of how information is gathered from the environment, not just how it is pooled after being discovered.

We can consider, in principle, what epistemic mechanisms might operate in groups. Steiner (1966) suggests potential group productivity (for a variety of group tasks, including both epistemic and physical tasks) is a function of three determinates: task demands, resources and processes. Task demands include the nature of the task itself, what sort of resources are needed, how much resources are needed and how the resources must be combined. The 'task demands' proposed by Steiner are equivalent to the 'epistemic systems' in Goldman's terminology. It is these institutional features which amount to social epistemic mechanisms. Agents' resources could include the intelligence and skill of individuals. The processes consist of the sets of actions taken by agents when they perform the task.

Steiner presents a taxonomy of five models of potential group productivity which are categorised according to the task demands: additive, disjunctive, conjunctive, compensatory and complementary models. In an *additive* model the task demands require each agent in a group to perform exactly the same actions and group performance will be the sum of the performance of individual agents. For example, a crew of stokers may be tasked with shovelling coal into a steam engine. The amount of coal shifted by the group is just the sum of coal shifted by each agent. Assuming no loss in productivity due to faulty processes (for example agents getting in each other's way) then as group size increases, group productivity increases. An additive model can also apply in epistemic settings. For example, a pub quiz team might be given a 'word-scrambler' puzzle where they are provided with a word such as 'heredity' and get points depending on how many other words they can form from the letters (such as 'here', 'red', 'tidy', and so on). As the size of a pub quiz team increases the combined knowledge of the team may increase and the number of points they get on this question may increase.

In a *disjunctive* model, group productivity is determined by the resources of the most able agent. For example, there may only be room for one agent to shovel coal into an engine. The maximum level of group productivity is limited to that of the strongest member of the group. In epistemic settings, a disjunctive model may also apply to a pub quiz team. For example, the group may face a question about the 2010 World Cup. The probability the group gets this question correct is limited to the competence of the group member who is supposed to be the expert on football.

In a *conjunctive* model the task demands require all agents in a group to perform a similar action. Group productivity is limited to the ability of the weakest agent in the group. For example, there may be several engines on a ship that must be fed coal at the same rate. The performance of the group of stokers as a whole is limited to the resources or ability (the strength and stamina) of the weakest member of the group. In epistemic settings, the pub quiz team may decide their answers to particular questions via a consensus. The probability of a correct consensus on a particular question is limited by the competence of the least competent member of the group.

The task demands of a *compensatory* model allow the actions of some agents in the group to offset the actions of other agents in the group. For example, some of the stokers may shovel coal at a rate so slow it risks starving the engine, while other stokers shovel coal so fast it risks suffocating the engine. On average the stokers shovel at just the right rate. If the size of the crew is small the engine won't receive coal at the correct rate. However as the size of the crew increases the slow stokers compensate for the fast stokers (and vice versa) and the engine receives coal at the correct rate. In epistemic settings, the pub quiz team may be asked to estimate the number of coins in a jar. On average agents will have a good idea how many coins there are. Although some agents may overestimate the number of coins in the jar and other agents may underestimate the number of coins in a jar, as the size of the group increases these under- and overestimates balance each other out and the group will tend to make the correct estimate.

Finally, the task demands in a *complementary* model can be divided and conducted by different agents. For example, shovelling coal may require both tall stokers, who can shift coal from the tender, and short stokers who can throw coal to the back of the fire box. A tall stoker would not be able to feed the fire properly by themselves. A short stoker would not be able to maintain the supply of coal by themselves. In epistemic settings, the pub quiz group may be asked how many wives Henry VIII had. One member of the group might think Henry VIII had eight partners. A different member of the group may know that Henry VIII had two mistresses. Between them these two agents should be able to deduce that Henry VIII had six wives.

Later in this thesis I will indicate which of Steiner's models apply to the various social epistemic mechanisms in political settings.

The processes by which groups of political agents can track the truth are summarised in the figure below. The figure shows how the most basic epistemic elements of information and agents are transformed via search and aggregation procedures into a collective judgement that tracks the truth. I also include in the figure an indication of which chapters in the thesis cover which parts of the overall process of group truth-tracking.

Figure 1.1: a summary of the two-staged process of group truth-tracking.

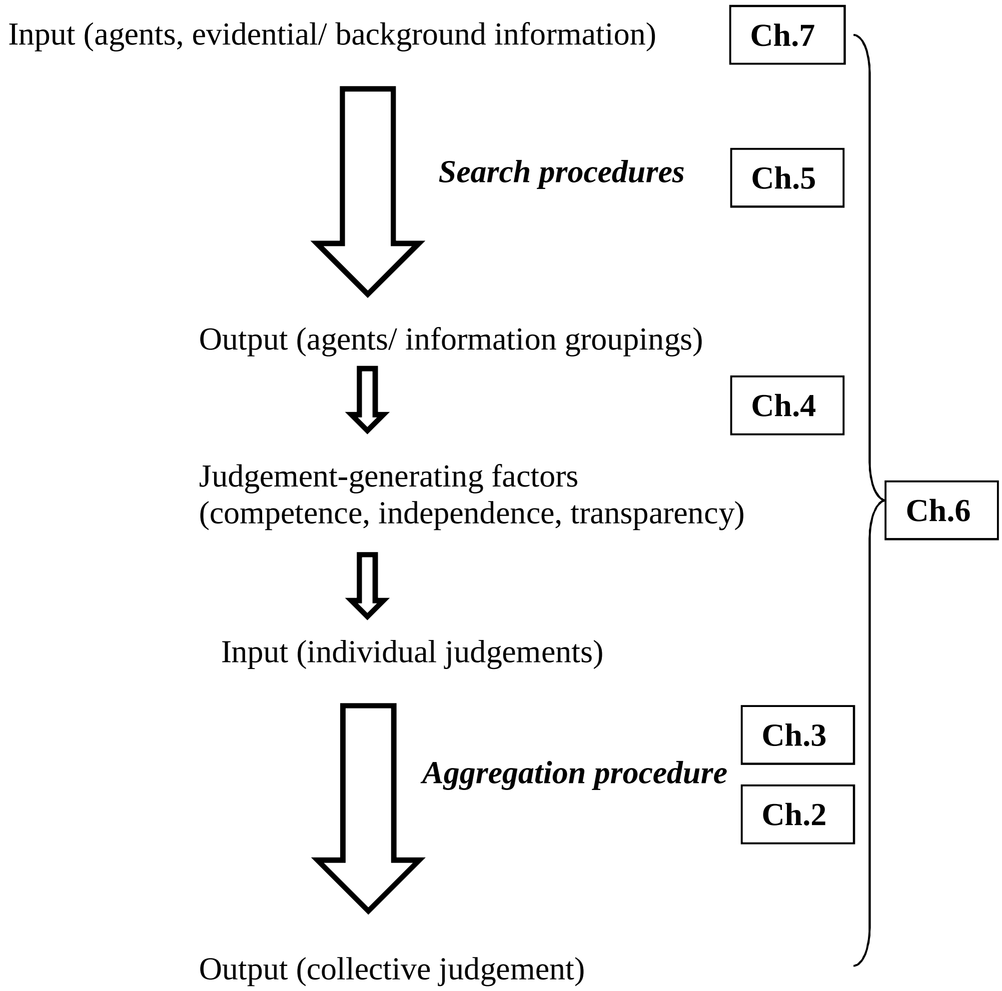

As will be argued, the key to truth-tracking by groups of political agents is, first, the identification of truth-conducive information by agents and, second, the aggregation of that information into the social choice. Institutional features in each of these stages amount to social epistemic mechanisms. However the thesis begins the presentation in reverse order. This is for two reasons. Firstly, the current literature on epistemic democracy often focuses on the aggregation procedure of majority voting, so rhetorically it makes sense to begin here. Secondly, understanding the inputs required for an aggregation procedure to track the truth will help in the analysis of the search procedures which are required to generate these inputs.

I will briefly summarise below the main points from the remainder of the chapters.

Aggregation procedures

A judgement aggregation procedure allows a group to generate a collective judgement (or social choice) based 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).

There are a variety of different aggregation procedures including (but not limited to) dictatorship, unanimity rule and majority rule. Under the aggregation procedure of dictatorship, an alternative will be the social choice if and only if a specific individual (the dictator) votes in favour of it. Under the aggregation procedure of unanimity rule, an alternative will be the social choice if and only if all the agents in the group vote in favour of it. Under the aggregation rule of majority rule, an alternative will be the social choice if and only if strictly more agents vote in favour of it than vote against it.

The epistemic performance of each of these three aggregation procedures is a function of the judgement-generating factors. There are four judgement-generating factors which are of interest:

- individuals' competencies, and the distribution of competencies in the group- the probability that agents will vote for the correct alternative;
- the transparency of competence- whether agents in the group or an observer can see the competencies of agents;
- the independence of agents- the probability that an agent will vote for the correct alternative, given the votes of other agents. If agents share information then they are more likely to vote in the same way; and
- group size.

The institutional features of each of the three aggregation procedures amount to social epistemic mechanisms. Given certain levels of competence, transparency of competence and independence relations, as group size increases the institutional features of the aggregation procedures make it more and more likely that the group will track the truth.

Dictatorship can meet the standards of baseline and absolute epistemic performance, provided that there is at least one individual agent in the group with high competence, whose competence is transparent. If the competencies of group members are heterogeneous then increasing group size is epistemically virtuous as it increases the probability that the group will contain such a high-competence individual.

Unanimity rule is a reliable aggregation procedure provided that we are only interested in avoiding an incorrect alternative as the social choice. The judgement-generating factors required for the aggregation procedure of unanimity rule to avoid the incorrect alternative as the social choice are a large number of agents whose levels of competence are greater than zero and who are conditionally independent. It does not matter whether the levels of agent's competencies are transparent or not. An alternative will only be the unanimity winner if every single agent votes for it. As the number of agents increases, the probability that every single one of the agents will vote for the same alternative decreases. Therefore, as group size increases, the probability that an incorrect alternative will receive a unanimous verdict decreases.

Majority rule can meet the three standards of epistemic performance if the competence of agents is better than random, if the distribution of competencies is symmetric about the mean and agents are independent. As group size increases the epistemic performance of the group improves. These claims are supported by the Condorcet Jury Theorem. In its classic form the Condorcet Jury Theorem states that if agents are 'competent' (the probability of agents voting for the correct alternative is homogeneous and greater than ½) and agents are 'independent', then the probability of a correct majority winner is monotonically increasing in group size and in the limit reaches certainty.

The institutional features of majority rule also amount to a social epistemic mechanism. If the probability of an agent voting for the correct alternative is greater than ½ there may still be a significant probability that this agent will vote for the incorrect alternative. If there is only a single agent or a small number of agents in the group then there may be a significant probability that a majority of them will vote for the incorrect alternative. However if the group size is large, the probability that a majority of the group will vote for the incorrect alternative will be small. The incorrect votes of the minority are offset by the correct votes of the majority.

The Condorcet Jury Theorem – agenda size and competence

The discussion in the thesis thus far will have assumed that agents are presented with an agenda comprising two alternatives, one correct and one incorrect. An obvious concern with this simplification is that in many real-world social choice problems there will be more than two alternatives. Multiple alternatives pose problems for the level of competence of agents. Whereas an individual might be quite competent at identifying the best alternative from a set of two alternatives, they may have more difficulty at selecting the best alternative out of a set of 100, 1000 or 100,000 alternatives. Arguably as the number of alternatives tends towards infinity, the competence of agents (the probability that they will vote for the best alternative) tends towards zero.

Increased agenda size poses particular problems for the aggregation procedure of majority rule. The classic Condorcet Jury Theorem states that if the level of competence drops below a half, the probability of a correct majority verdict *decreases* as group size increases and in the limit tends to *zero*. The classic

Condorcet Jury Theorem cannot cope with agendas comprising more than two alternatives.

I consider two main extensions to the classic Condorcet Jury Theorem to cope with multiple alternatives. Firstly, I consider Condorcet's own extension (as presented in Young (1988)) which requires a pair-wise comparison between each of the alternatives. I conclude that, when implemented, this extension requires too much effort on behalf of agents. I also consider the extension of List and Goodin (2001) which extends the classic Condorcet Jury Theorem from majority voting on a two-placed agenda to plurality voting on a many-placed agenda. I conclude that the application of the List and Goodin extension suffers from the same problem discussed above, namely, that as group size increases the competence of agents will decrease. The low level of agent competence may mean the probability of a correct plurality winner is too low.

Using the insights from both the Condorcetian and List and Goodin extensions of the classic Condorcet Jury Theorem I argue for a mixed approach for coping with multiple alternatives. A social planner can use multiple elections with agendas of varying sizes and groups of varying sizes to balance the competing demands of reducing the burden on voters (by minimising the number of elections they participate in) and increasing the competence of agents (by reducing the size of the agendas they face).

I also address the 'Disjunction Problem', as presented in Estlund (2008). The Disjunction Problem makes use of the List and Goodin (2001) extension of the

Condorcet Jury Theorem to multiple alternatives to challenge the fulfilment of the competence assumption. The crux of the Disjunction Problem is that there is no principled way to determine the number of alternatives that should be on an agenda. If we cannot determine the number of alternatives that should be on an agenda, we cannot determine the level of competence required for the competence assumption of the Condorcet Jury Theorem to hold. I clarify the Disjunction Problem and argue that what it actually shows is that the framing of an agenda by a social planner can determine whether the competence assumption of the extended Condorcet Jury Theorem does or does not hold. There is no way in principle to ensure that the agenda will be set in such a way that the competence assumption does hold. However I argue that any attempt to justify the competence assumption 'in principle' is misguided. There is always a possibility a social choice problem will include misleading information. The best hope of defending the competence assumption is identifying an appropriate reference class of social choice problems where the competence assumption is likely to hold. Identifying a suitable reference class of problems is not something that can be done analytically.

The generation of the inputs to aggregation procedures

The existing accounts of epistemic democracy that focus on aggregation procedures only give conditional support to the truth-tracking ability of groups. They show how groups can track the truth given certain types of judgement-generating factors. They are silent on how these judgement-generating factors are themselves generated or whether they are plausible. More particularly, the existing aggregative accounts of group truth-tracking begin at the point at which agents already have a set level of competence, in a particular distribution, with certain independence relations holding, and the transparency or otherwise of competence pre-determined. But it cannot be taken as given that agents will have information regarding the correct alternative on an agenda. Nor can it be taken as given that the required independence relations will hold or that the transparency of competence is established. We need an account for how the features of a group of agents, including competence levels, transparency of competence and independence relations, develop. The truth-tracking institutional features of some aggregation procedures can provide a conditional epistemic justification for group decision making; an account of the formation of the judgement-generating factors will provide the antecedent to this conditional justification.

The competence of an agent is defined as the probability that this individual agent votes for an alternative, given that it is correct. The competence of an agent represents the probability of an event occurring, namely the probability that a particular agent will vote for the correct alternative. The agent's vote for a particular alternative is determined by the combination of their causal influences. I utilise the taxonomy of causal factors presented Dietrich (2008). The causal factors determining an agents vote (and therefore the probability that the agent will vote for the correct alternative) can be truth-conducive or they can be misleading. Truth-conducive causal factors will make an agent more likely to vote for the correct alternative; misleading factors will make an agent less likely to vote for the correct alternative. Causal factors can be evidential or background. Evidential factors are causal relatives of the true state of the world. Background factors are not causal relatives of the state of the world, but nevertheless allow an agent to interpret evidential information.

Causal factors (be they evidential or background, truth-conducive or misleading) can either be held privately by agents or held in common between agents. If all causal factors of agents' votes are held privately then agents will be independent, conditional on the state of the world. If however agents have at least some evidential or background factors in common, there will be certain dependence relations in the votes of agents.

Finally, the nature of the causal factors determining an agent's vote will also determine whether an agent's competence is transparent or not. For example, if the evidence generating an agent's competence is of a kind that can be shown to other agents, her competence will be transparent to other agents.

At this point in the thesis I leave consideration of aggregation procedures and move on to search procedures.

Group search procedures

I provide a general framework for search procedures involving groups of agents. A single agent searching for an object of interest may only have a small probability of finding it. But if we employ a group to search for the object the probability that at least one of the group members will find it can be significantly higher. I present a theorem that states under certain assumptions the probability that a group of agents will identify a particular object is increasing in group size and in the limit tends to certainty. There are two different mechanisms behind the epistemic performance of a group search procedure. Firstly, increasing the number of agents in the group can increase the probability that an agent will visit the location of the object. Secondly, if we increase the number of agents visiting the *same* location we can increase the probability that the object at a particular location will be recognised by a member of the group. The assumptions of the theorem are modified to produce extensions of the theorem.

I also develop a model of a group search procedure to investigate the dynamics of group search. In the model there is a set of locations, one of which contains the object of interest. Individual agents engage in a search for the object by moving from location to location. The locations an agent visits are determined by four agent-specific variables: their initial partitioning of the search space, the convention the agent employs for ordering the locations, the start point of their search and their search heuristic. The objects an agent finds are determined by the locations they visit and their capacity to recognise objects at those locations. If there are differences in the locations visited by agents and / or differences in the ability of agents to recognise objects then as group size increases the probability that a member of the group finds the object increases and in the limit reaches certainty.

The model of the group search is reproduced in the computer program 'NetLogo' and subjected to simulations. The results of the simulations confirm both the formal results of the search theorem and the conceptual arguments of the search model: as group size increases the probability an object will be found tends to certainty and is increasing up to the limit. The simulation results also show the impact on a group's search performance of adjusting the agent-specific search variables.

On its own the institutional features of a group search procedure, as presented in the model and backed by both the search theorem and the simulation results, amount to mechanisms by which groups of agents can track the truth. In addition a group search procedure can link in with the aggregation procedures described earlier. In the subsequent chapter I will explain how search procedures can be used to fill in some gaps in accounts of epistemic democracy which rely on aggregation procedures.

The link between search procedures and aggregation procedures

I claim that truth-tracking by groups of political agents occurs via two procedures. Standard epistemic defences of democracy often focus on aggregation procedures such as majority rule, which pool the information individual agents have regarding the true state of the world. I also put the case for groups of agents employing search procedures to find information in the first place. The institutional features of search procedures and aggregation procedures amount to social epistemic mechanisms.

I provide an account of how the search and aggregation procedures link up. Search procedures allow groups of agents to extract information from the environment. Aggregation procedures allow individual agents within a group to share the information they have with the wider group.

The linking of search procedures to aggregation procedures fills two gaps in current epistemic defences of democracy that rely on aggregation procedures. Firstly, current accounts of aggregation procedures specify the types of judgement-generating factors (competence, independence, and transparency) required for a group to track the truth but they are silent on how the judgement-generating factors form. Search procedures can be used by agents to search for evidential and background information to develop their levels of competence. And diversity in the search procedures of individual agents will generate the dependence relations in the group. Secondly, search procedures can be employed by a group to find possible alternatives and to set the agenda for a social choice.

This chapter also gives consideration to nested social choice problems. Any social choice in fact involves two procedures (a search procedure followed by an aggregation procedure) and there are epistemic advantages to increasing group size in each procedure. But the final social choice, the alternative which the group judges to be the true state of the world, may in fact be the result of a sequence of different (two-staged) social choices including a choice over the topic to consider, a choice over how to assess the quality of alternatives, a choice over the alternatives to place on the agenda, and finally a choice over the alternative to be the social choice.

Once our framework for group truth-tracking joins search and aggregation procedures together we can consider the interaction between the two. We can see how contingencies in the way a search procedure is conducted mean particular aggregation procedures will be optimal at tracking the truth. Similarly, if an institutional decision is made in advance to use a particular aggregation procedure then this will influence the way in which a search procedure should be conducted so that it generates the appropriate levels and distributions of competencies and independence relations.

The limits of the informational environment

The final substantive chapter focuses on contingencies in the informational environment which place restrictions on the absolute epistemic performance of aggregation procedures. These issues are discussed mainly via the framing of majority voting and the Condorcet Jury Theorem. This is because much of the relevant literature focuses on the asymptotic limit of the Condorcet Jury Theorem.

Firstly I address the problem of the possibility of misleading information. A small but significant literature on this topic has developed quite recently. The analysis in this literature shows that the mere possibility of misleading evidence and background factors means the asymptotic limit of the Condorcet Jury

Theorem is not certainty but some value less than certainty. This means that the absolute epistemic performance of majority voting may be too low: groups using majority voting as an aggregation procedure may not be very reliable at identifying the true state of the world.

Secondly, I address the problem of finite information. In some social choice problems there may simply be insufficient information for a group to determine the true state of the world, no matter what search or aggregation procedures the group employs. I consider what institutional responses a group might employ to maximise the probability of a correct social choice when the amount of truth-conducive information is limited. If information is finite, the truth-conducive value of the information will be maximised by agents sharing the information. Agents can share truth-conducive information and increase their levels of competence. The agents will remain independent (conditional on common factors) provided they have at least some background factors held privately.

Conclusion

Understanding the mechanisms by which groups of political agents can track the truth has obvious normative implications. If it is the case that some social choice problems have a correct alternatives then understanding the conditions under which a group of agents can identify this alternative allows us to implement the appropriate institutional arrangements.

I argue that there are two main steps or procedures operating in groups of political agents. There is a search procedure by which agents identify truth-conducive information in the environment. Subsequently there are aggregation procedures which pool this truth-conducive information. The institutional features of both search and aggregation procedures account for the ability of a group to track the truth and amount to social epistemic mechanisms. Increasing group size is an important feature of all the social epistemic mechanisms and this fact lends support to the epistemic importance of including a large and diverse a group of agents in political decision making. However the two-staged framework of search then aggregation, and the social epistemic mechanisms operating in each stage, do not necessarily support a strictly democratic form of decision making. For example, while majority or plurality rule may be the epistemically optimal aggregation procedure in some social choice problems, there will be other social choice problems where expert dictatorship is optimal.

The sharing of information, after the conclusion of the search procedure but before the aggregation procedure, can be epistemically virtuous. Building an appropriate model of deliberation to fit within the framework of search and aggregation is a topic set aside for future research.

Summary of notation

Here I summarise the main pieces of notation that I will employ in the remainder of the thesis.

P = the probability of a correct social choice.

P^+ = the positive reliability of a group, the probability the group chooses the alternative given that it is correct.

P^- = the negative reliability of a group, the probability the group will avoid an alternative given that it is incorrect.

$i, j, k, \ldots$ = variables for individual agents.

1,2,3,…n = names for individual agents.

x = the state of the world. x can take two values: 0 or 1.

p = the homogeneous level of competence of agents, the probability that agents will vote for the correct alternative.

p_i = the competence level of some unknown agent i.

p_1 = the competence of agent 1.

p' = the average competence of a group of agents.

p_i^0 = the prior competence of agent i, their level of competence before they receive any evidential information.

p_i^1 = the posterior competence of agent i, their level of competence at time t_1 after they have received some evidential information.

n = the number of agents in a group.

$k_1, k_2, k_3 \ldots k$ = the names for alternatives.

v_i = the vote of agent i.

e = the number of elections.

c_1^b = a background cause of an agent's judgement, in this case cause number 1.

c_2^e = an evidential cause of an agent's judgement, in this case cause number 2.

$\Pr(S_i)$ = the search competence of agent i, the probability that the agent moves to a particular location containing an object of interest.

$\Pr(R_i|S_i)$ = the recognition competence of agent i, the probability that the agent recognises a particular object at a location given that they move to that location.

P_S^n = the probability of a member of a group visiting the location of an object.

P_R^n = the probability of a member of a group recognising an object at a particular location.

P_F^n = the probability of a member of a group finding the object of interest.

I employ a convention for subheadings in the thesis where **first-level subheadings are in bold,** <u>second-level subheadings are underlined</u> *and third-level subheadings are in italics.*

Chapter 2: Aggregation procedures.

The core concern of is thesis is the mechanisms by which groups of agents can track the truth. Different agents may have different judgments as to the true state of the world. This chapter sets out the particularly salient procedures by which individual judgements can be aggregated into a collective judgement or social choice, namely dictatorship, unanimity rule and majority rule. The chapter establishes the probability that these aggregation procedures will generate a social choice that is correct and the probability they will avoid a social choice that is incorrect. There are four key judgement-generating factors which determine the probability a given aggregation procedure will track the truth: the competence of agents and distribution of competencies in the group; the transparency of agents' competencies; the independence of agents; and the group size. Given appropriate judgement-generating factors each of the aggregation procedures can successfully track the truth. Importantly, there are certain combinations of judgement-generating factors where increasing group size is epistemically virtuous. The analysis of the three aggregation procedures in this chapter provides an explanation for their truth-tracking ability, for how the institutional features of the aggregation procedures can operate as social epistemic mechanisms. This includes an explanation of how the classic Condorcet Jury Theorem works and the importance of its competence and independence assumptions.

The framework for aggregation procedures

Proponents of epistemic democracy argue that democratic forms of decision making are desirable in so far as they track the truth. For epistemic democracy to have any purchase it must be the case that at least some political decisions are judgements about matters of fact, about the actual state of the world. For example, whether a nation possesses a nuclear weapon 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 matters of fact. Propositions which describe possible states of the world are termed *'alternatives'*[4]. 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

[4] A policy choice may involve a series of choices over alternatives. For the sake of simplicity I assume that where there is a series of choices over propositions the propositions are not logically interconnected (unless otherwise stated). Where this does not hold, and there is a logical interconnection between propositions then we can run the risk of a discursive dilemma (see List, 2006).

the agenda only contains some of the possible alternatives then there is a risk that the correct alternative is not included.

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[5] allows a group to generate a collective judgement (or social choice) based 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)[6].

Figure 2.1: aggregation procedures.

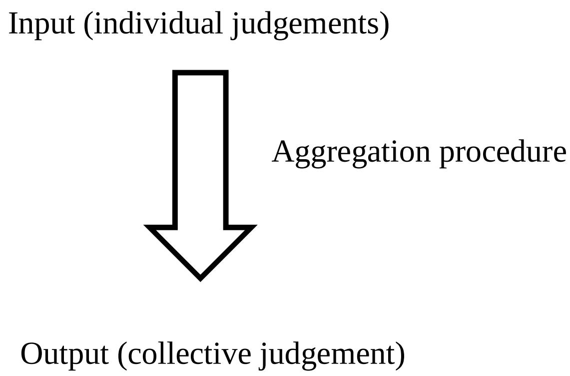

[5] In this thesis I am interested in judgement aggregation procedures, rather than aggregation procedures more generally.

[6] However not all aggregation procedures need be functions. A function requires that each input (or combination of votes) have a unique output (social choice), but there may be some aggregation procedures (perhaps including some deliberation) which could have a variety of outputs depending on contingencies in the way the inputs are treated. Figure 2.1 also comes from List (2008).

There are a variety of different aggregation procedures including (but not limited to) dictatorship, unanimity rule and majority rule[7]. 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[8]. Each aggregation procedure has different virtues, but the concern of this thesis is the epistemic virtue, the probability that the aggregation procedures will select the correct alternative (and avoid the wrong alternative) as the social choice[9].

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)[10].

To determine the epistemic performance of different aggregation procedures we need to, firstly, draw a distinction between positive and negative reliability. The

[7] If n represents the number of agents then the number of distinct possible aggregation procedures for a dichotomous choice is given by the formula 2^{2^n}. So, for example, if there are three agents in a group then there are $2^{2^3} = 256$ possible aggregation procedures (Christian List, unpublished lecture notes).

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

[9] As for the non-epistemic (or procedural) virtues of these aggregation procedures: dictatorship is the only aggregation procedure that meets the conditions of Arrow's theorem (universal domain, Pareto efficiency and independence of irrelevant alternatives); majority rule is the only aggregation procedure that meets the conditions of May's theorem (universal domain, anonymity, neutrality and positive responsiveness).

[10] 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. The issue of agendas with multiple alternatives is addressed in the next chapter.

positive reliability, P^+, is the probability that a group using a particular aggregation procedure will judge an alternative to be true given that it is true. The negative reliability, P^-, is the probability that a group using a particular aggregation procedure judges an alternative not to be true given that it is false. Because positive and negative reliabilities are probabilities, they take values in the interval [0,1]. Given this distinction between positive and negative reliabilities there are four possible judgements that can be made, as shown in the table below:

Figure 2.2: possible group judgements.

	Judgement: true	**Judgement: false**
State of the world: true	Positive reliability P^+	False negative $1 - P^+$
State of the world: false	False positive $1 - P^-$	Negative reliability: P^-

For both positive and negative reliability there are three measures of group epistemic performance that are of interest. Firstly, we have a measure of 'baseline epistemic performance' which is the probability that a group will be better than random at picking the true alternative (avoiding the false alternative). Given that we only have two possible alternatives then a group would have a 0.5 probability picking the correct alternative at random, so the measure of baseline epistemic performance requires P^+ or $P^- > 0.5$. Secondly, 'relative epistemic performance' is a measure of the epistemic performance of a group when compared to an individual member of that group. If we are concerned with the truth-tracking ability of groups then we need groups to be better than individual members of the group at identifying correct alternatives/ avoiding incorrect alternatives. Finally, 'absolute epistemic performance' is a measure of the probability the group will select the correct alternative as the social choice. This takes a value in the interval [0,1], and we would need P^{+} or P^{-} to be very close to 1 (very likely to select the correct alternative/ avoid the incorrect alternative) if we want to point to a group as being a successful truth-tracker.

We can summarise the six ways in which we can assess the epistemic performance of groups using the various aggregation procedures:

<u>Positive reliability</u>

Baseline	A group is better than random at selecting the correct alternative
Relative	A group is better than an individual at selecting the correct alternative
Absolute	A group is good at selecting the correct alternative

<u>Negative reliability</u>

Baseline	A group is better than random at avoiding the wrong alternative
Relative	A group is better than an individual at avoiding the wrong alternative
Absolute	A group is good at avoiding the wrong alternative

The positive and negative reliabilities of a group, and the group's baseline, relative and absolute measures of epistemic performance, depend crucially on both the aggregation procedure the group employs and on the judgement-generating factors to the aggregation procedure. The inputs to an aggregation procedure are the judgements or votes of individual agents. On a two-placed agenda the votes of agents are typically recorded as a 1 if an agent votes for the first alternative on the agenda and a 0 if the agent votes for the other alternative on the agenda (which is equivalent to not voting for the first alternative). In an epistemic setting the votes of agents for or against an alternative will be determined by a combination of causes, including the truth-conducive information that an agent has received. A social planner or observer will typically be unaware of all the causes of agent's votes and as such does not know in advance whether a particular agent will vote 1 or 0 and whether an aggregation procedure will generate the correct social choice. Instead the social planner may be aware of certain causal factors which generate the inputs to the aggregation procedure and can attribute a probability to the event of an agent voting correctly. The way in which these inputs or judgements are generated can be classified according to a taxonomy that focuses on four variables:

- individual agents' *competencies* and the distribution of competencies in the group: the probability that each individual agent will vote for the correct alternative, given the state of the world;

- the *transparency* of competence: whether agents in the group can see the competence of other agents (or whether an observer or social planner can see the competence of agents);
- the *independence* of agents: the probability that an agent will vote for the correct alternative, given the votes of other agents and the state of the world; and
- *group size*.

We will consider the judgement-generating factors to an aggregation procedure first, before going on to consider the aggregation procedures themselves.

Judgement-generating factors

<u>Competence</u>

Individual agents will have a positive reliability and a negative reliability, just as the group does. The positive reliability of an agent is the probability that the agent will judge an alternative to be true, given that it is true. The negative reliability of an agent is the probability that the agent will not judge an alternative to be true, given that it is false. In many cases the positive and negative reliabilities will be identical, but in some cases they may be different. For example, consider two types of non-human agents. A scanner at airport security might have a high positive reliability but lower negative reliability; it might have a high probability of registering 'true', given the presence of a metal object on a passenger and a lower probability of registering 'false' given that a passenger does not have any metal objects on them. The airport scanner is designed to have this asymmetry between positive and negative reliability because the costs of a false negative (letting a knife onto a plane) are high but the costs of a false positive (having to 'pat down' a passenger) are low. Similarly it may also be possible to have high negative reliability and lower positive reliability. For example, a test for blood alcohol may have a high probability of registering 'false' if it is false that a suspect has alcohol in his blood stream, but a lower probability of correctly registering 'true' given that a suspect does have alcohol in his blood stream. In this example it may be judged that wrongly convicting a motorist of drunk-driving is worse than not convicting a drunk-driver.

For the sake of simplicity we will assume that the positive and negative reliabilities of an agent have the same value; an agent is equally able to correctly judge a proposition is false, given that it is false, as they are to correctly judge a proposition is true, given that it is true. Each agent i has a level of competence p_i , which is the probability that they will vote for the correct alternative. Because the competence of agents is a probability it takes a value in the interval [0,1]. Under this simplification there are two possible judgements an individual agents can make:

Figure 2.3: possible individual judgements, given identical positive and negative reliabilities.

	Judgement: X	**Judgement:¬X**
World: X	Correct: p_i	Incorrect: $1 - p_i$

The distribution of competencies in a group depends on the competencies of individual agents in the group. So, for example, if our group is comprised of agents $1, 2, \ldots, n$ then the average competence of this group depends on the competence of agent 1 and agent 2 … and agent n.

Transparency of competence

To say that competence is transparent is to say that all agents (or an observer or social planner) know the competence levels of all the agents in the group and they know that they know the competencies. To say that competence is opaque is to say that agents (or an observer or social planner) do not know the competencies of all the agents, and they know that they do not know the competencies of the agents. This treatment of transparency involves three important assumptions. Firstly, it is assumed that transparency is a binary notion. Competence is either transparent or it is opaque. Secondly, we assume that the transparency or opaqueness of competence is homogeneous across a population and is determined by the contingent circumstances of a particular social choice problem. Finally we ignore cases where agents (or an observer or social planner) lack self-awareness of their knowledge of competence. So we ignore cases where agents don't know they don't know competencies and we ignore cases of 'blind-sight' where agents do actually know the competencies, but are not aware that they know the competencies. All three of these simplifying assumptions can be legitimately challenged. There may be degrees of transparency, and these may be heterogeneous across the population[11].

[11] Heterogeneous transparencies would raise interesting questions of how to attribute weights to different agents. Is it better to trust an agent who you are certain has a competence of 0.6, or take a chance on an agent you are 0.8 confident has a competence of 0.9?

Furthermore there could be situations in which agents are not aware they do not know competencies and this could create the potential for errors[12]. However, we are primarily concerned with cases in which groups of agents can track the truth. Cases where agents know that they know the competencies and cases where agents know they do not know the competencies are the two types of social choice problem that are directly relevant to the truth-tracking ability of the aggregation procedures below.

Independence of agents

Agents are independent if the probability of them voting for an alternative, given the state of the world, is identical to the probability of them voting for an alternative, given the state of the world AND the vote of another agent. The variable of independence captures the extent to which there is diversity in the voting behaviour of agents in a group. At one extreme, where all agents are independent conditional on the state of the world, the vote of one agent tells us nothing about how another agent will vote. At another extreme, where agents are entirely dependent, all agents vote identically and so the vote of one agent will tell us precisely how all other agents will vote. If full-blown independence is violated, if the probability of an agent voting for an alternative given the state of the world is *not* identical to the probability of them voting for an alternative given the state of the world and the vote of another agent, then independence might be secured by conditionalising on the factors held in common between agents. So agents will be conditionally independent if the probability of an

---

[12] For example, agents might have the certain (but entirely mistaken) belief that agent i has competence of 1.0, and as a consequence make the wrong agent dictator.

agent voting for an alternative given the state of the world and any factors held in common between agents is identical to the probability of them voting for an alternative given the state of the world, any factors held in common between agents and the vote of another agent.

<u>Group size</u>

Group size, the number of agents who are permitted to express a judgement on an agenda, can have a significant impact on the probability that a certain aggregation procedure will deliver the correct social choice. Group size is represented formally as n.

With the taxonomy for the judgement-generating factors of the aggregation procedures now set out, we can move on to consider what combinations of judgement-generating factors are required for the different aggregation procedures to track the truth as group size increases.

Aggregation procedures

<u>Dictatorship</u>

Under the aggregation procedure of dictatorship, the social choice is determined by one individual. The positive reliability of the group is therefore identical to that of the dictator:

$$P^{+} = p_i$$

With dictatorship the group's negative reliability is always identical to its positive reliability in virtue of the assumption that an agent’s positive reliability is identical to the agent’s negative reliability, for all agents (including the dictator).

If the competence of agents is *transparent* then the epistemically best the group can do is if the most competent member of the group is made the dictator. The epistemically worst the group can do is if the least competent member of the group is made the dictator. If the competence of agents, including the dictator, is not transparent then the probability of a correct social choice may simply be unknown.

Violations of *independence* are not epistemically disadvantageous for the aggregation procedure of dictatorship. Only a single agent gets to cast a vote and so the conditional probability of an agent voting correctly given the vote of another is irrelevant. In fact if agent’s votes are determined by the truth-conducive information they receive it is epistemically beneficial to share this information and violate independence, because in sharing information the competence of agents (including the dictator) increases.

For the aggregation procedure of dictatorship to track the truth it is beneficial for individual agents to be as *competent* as possible. This increases the probability that the agent selected at random from the group will be of high competence. And if competence is transparent and we are able to select the most competent agent to be dictator, then maximising the competence of the dictator is obviously of value.

If, as is plausible, the competence of a population of agents is heterogeneous and ranges between 0 and 1 then as *group size* increases the probability of a correct social choice can also increase. This is because increasing group size increases the probability that the group will include individuals with high competence. Suppose we form a group of agents by taking samples from a wider population with heterogeneous competencies ranging from 0 to 1. We can define the event of sampling an agent with the maximum level of competence as D_i. We can assume that the probability of sampling any one agent who has a level of competence at the maximum level is independent of the event of sampling another agent who has a level of competence at the maximum level. For each group of size *n*, $\bigcup_{i=1}^{n} D_i \subseteq \bigcup_{i=1}^{n+1} D_i$, and hence by the monotonicity of probability $\Pr(\bigcup_{i=1}^{n} D_i) \leq \Pr(\bigcup_{i=1}^{n+1} D_i)$.

If the competence of agents' is transparent then it is possible, ceteris paribus, to make a high competence agent the dictator. In such cases, the relative and absolute epistemic performance of dictatorship as an aggregation procedure is good. If the group contains at least one agent of high competence then making this agent dictator means the social choice chosen by this person is likely to be correct and more likely to be correct than that chosen by any of the other agents. Furthermore the baseline epistemic performance of the group using dictatorship is good since if the dictator has high competence they are more likely to select the correct alternative than a random choice.

If however the competence of agents is not transparent, then increasing group size may increase the probability that the competence of the dictator is the expected value of competence[13]. However a dictator with this expected value of competence may not exist[14]. Importantly, as group size increases the competence of the dictator (and therefore probability of a correct social choice) does not increase. Furthermore, when competence is not transparent, the relative and absolute epistemic performance of the group can be poor (although it may be better than baseline reliability if average competence is greater than 0.5).

Henceforth I will term the aggregation procedure of dictatorship ‘expert dictatorship’, since I am interested in aggregation procedures that can track the truth and dictatorship only tracks the truth when the dictator is an expert. The institutional features of expert dictatorship, namely the stipulation that the judgement of one agent will determine the social choice and that the most competent agent will be selected for this role, amount to a social epistemic mechanism. Increased group size is epistemically virtuous for expert dictatorship since increasing group size tends to increase the level of competence of the expert dictator. Under the framework of Steiner (1966), discussed briefly in the introduction, expert dictatorship is a disjunctive model of group productivity. Only a single agent from the group performs the group’s task (selects the social choice) and the performance of the group (the probability

[13] In accordance with the law of large numbers.

[14] For example, if half of all agents have a competence of 0.9 and half of all agents have a competence of 0.5, then the expected value of competence is 0.7. However an agent with competence of 0.7 does not exist.

of a correct social choice) is limited to the ability (the competence) of that single agent.

In a political setting, a form of expert dictatorship is employed where decisions are delegated to a Government Minister. There may be far too many day to day decisions in government for the Cabinet to consider as a group. Decisions in the defence portfolio will be the responsibility of the Minister of Defence. The Minister of Defence will be provided with detailed briefings by his or her Ministry. In addition, if the Prime Minister wants correct decisions to be made in this portfolio, he or she will appoint a Minister who has sufficient capability to make correct decisions. As such, we should expect the Minister of Defence to have a high level of competence; we should expect him or her to be an expert. As the size of a Government's majority in Parliament increases, the 'pool of talent' should also increase; the probability that the Government will include a member of exceptional ability, whom the Prime Minister can appoint as Minister of Defence, increases.

Unanimity rule

Under the aggregation procedure of unanimity rule, an alternative is the social choice if and only if every individual in the group votes for it. The positive reliability of the unanimity rule, the probability that the group will select the correct alternative as the social choice, assuming the votes of agents are independent is:

$$P^{+} = \prod p_i \leq p_i$$

The higher the *competence* of individual agents, the more likely it is that the aggregation procedure of unanimity rule will select the correct alternative as the social choice. The epistemic performance of the group is limited by the least competent member of the group therefore it is epistemically best if agents have as much truth-conducive information as possible. It is better that the truth-conducive information is spread around evenly in the group rather than being concentrated in the hands of just some of the agents.

The *transparency* of competence does not have any impact on the actual ability of the group of agents to track the truth. However for us to know the probability of a correct social choice we do need to know what the competencies of the different agents are.

Unless the competence of agents is 1.0, increasing *group size* decreases the probability the group will identify the correct alternative i.e. increasing group size is epistemically harmful to the positive reliability of unanimity rule. This is because the probability of a series of events occurring is less than or equal to the probability of the individual events occurring. A unanimous verdict for the correct alternative requires all the agents in a group to vote for the correct alternative. Suppose we start with a group size of one, which just includes agent i. The probability of a unanimous verdict for the correct alternative is:

$$P^{+} = p_i$$

Now we increase our group so that it is comprised of agents i and j. For there to be a unanimous verdict for the correct alternative now, both agent i and agent j need to vote for the correct alternative. The probability for this occurring is:

$$P^+ = p_i \times p_j$$

Furthermore:

$$p_i \times p_j \leq p_i$$

Unless agent j is guaranteed to vote for the correct alternative (unless $p_j = 1.0$) then adding j to the group decreases the probability of a unanimous verdict for the correct alternative.

Violations of *independence* (due to the sharing of information) are potentially epistemically advantageous for positive reliability unanimity. Suppose the competence of agents conditional only on the state of the world (p) is homogeneous and $p = 0.6$. If agents are probabilistically independent given the state of the world then the probability of a correct social choice is:

$$P^+ = 0.6^n$$

As $n \to \infty,\ P^+ \to 0$. If agents are probabilistically independent given the state of the world, then as stated above as group size increases towards infinity the probability of a correct social choice tends towards zero. But if agents have identical information that determines their votes, then the conditional probability of all $n - 1$ agents voting correctly given state of the world and given that the first agent votes correctly is 1.0. As such, where independence is violated:

$$P^{+} = 0.6 \times 1.0^{n-1} = 0.6$$

Where agents are not conditionally independent (when they are probabilistically dependent) the probability of a correct social choice is identical to the probability that a single individual agent will select the correct alternative. As such, increasing group size makes no difference to the probability the group will select the correct alternative.

For positive reliability P^{+} (the probability of selecting the correct alternative as the social choice) the aggregation procedure of unanimity rule fails the tests of baseline, relative and absolute epistemic performance. If we assume that the competence of agents is less than 1.0 and agents are independent then as group size increases the probability of a correct unanimous social choice tends towards zero. As such the group will tend to be less likely than a random choice at selecting the correct alternative and the group will be less likely than an individual to select the correct social choice. Furthermore the group will be unlikely simpliciter to select the correct social choice.

Although unanimity rule is poor in terms of positive reliability, it does well in terms of negative reliability P^{-} (the probability of avoiding an incorrect alternative as the social choice). If p_i is the probability that an agent will vote for the correct alternative, then $1 - p_i$ is the probability that an agent will vote for the incorrect alternative. If we assume the competence of agents is homogeneous then the probability that there will be a unanimous vote in favour of the incorrect alternative is $(1 - p_i)^n$. Therefore the probability of a group using unanimity avoiding the incorrect alternative as the social choice is:

$$P^- = 1 - (1 - p_i)^n$$

Provided that agents are not totally incompetent (provided that p_i is not zero) then as $n \to \infty, P^- \to 1$ i.e. the probability of not selecting the incorrect alternative as the social choice tends towards certainty as the group size tends towards infinity. The mechanism that drives the good epistemic performance of negative reliability unanimity is similar to the mechanism that drives the poor epistemic performance of positive reliability unanimity. A unanimous verdict for the incorrect alternative requires all the agents in a group to vote for the incorrect alternative. The probability of a series of events occurring is less than or equal to the probability of the individual events occurring, so as group size increases the probability of a unanimous verdict decreases towards zero. 'A unanimous verdict for the incorrect alternative/ not- a unanimous verdict for the incorrect alternative' is a complete logical partition and so the probability of one of these events occurring is certainty. If the probability a unanimous verdict for the incorrect alternative tends towards zero as group size increases, the probability of *not* having a unanimous verdict for the incorrect alternative tends towards certainty.

The ideal judgement-generating factors for unanimity rule to be successful at negative reliability are high *competence*, *independence* and (importantly) large *group size*. *Transparency* of competence is not important for avoiding incorrect alternatives (though it is important for knowing how likely the group is to avoid incorrect alternatives). Independence is important because if the conditional probability of one agent voting correctly given the state of the world and the fact that another agent votes correctly is 1.0, then the probability of the group avoiding the incorrect alternative as the social choice would be identical to the competence of an individual agent.

In terms of negative reliability, the baseline, relative and absolute epistemic performance of unanimity rule as an aggregation procedure is good. As group size *n* increases, unanimity rule will be better than random, better than an individual and likely to avoid the wrong alternative as the social choice. However there is a cost associated with negative reliability unanimity rule, namely that there is a high probability of no social choice.

Henceforth we will term the aggregation procedure of unanimity rule 'negative reliability unanimity rule' since we are interested in the aggregation procedures that can track the truth.

According to the Steiner (1966) taxonomy, the institutional features of unanimity rule would be a conjunctive model of group productivity since all agents have to perform the same action to get the desired outcome (all agents have to vote for the correct alternative for the correct alternative to be the social choice). The institutional features of the aggregation procedure of negative reliability unanimity rule are a social epistemic mechanism and would fit under the category of a disjunctive model of group productivity, since it only takes a single agent to vote for the correct alternative for the group to avoid the incorrect alternative as the social choice.

Negative Reliability Unanimity Rule may be employed on criminal jury trials. Here it is thought the consequences of a false negative are better than the consequences of a false positive; that it is better to let a guilty person go free than to wrongly convict an innocent person. Requiring a jury to have a unanimous verdict, and increasing the size of a jury from a single judge to twelve jurors, increases the probability of avoiding convicting the innocent. In a political setting, negative reliability unanimity rule might be employed in cases where the consequences of a bad status quo policy are better than the consequences of moving to a new incorrect policy. For example, pre-emptively attacking Iran on the mistaken assumption that they have nuclear weapons may be worse than forgoing the opportunity to attack Iran if indeed they do have nuclear weapons. By insisting on a consensus on the decision to attack, and by including all twenty-three members of Cabinet rather just the Minister of Defence alone, the Government would increase the probability of avoiding a disastrous policy choice.

Majority rule

As with dictatorship, the group's positive and negative reliabilities are identical under majority rule[15]. Under majority rule, the social choice is the alternative that more than half of the individual members of a group vote for. Under majority rule, the probability that the group selects the correct alternative as the social choice is given by[16]:

$$P = \sum_{S \subset N : |S| > \frac{n}{2}} \prod_{i \in S} p_i \prod_{i \notin S} (1 - p_i)$$

If the competence of agents in a group is heterogeneous and symmetric about the mean then the following formula gives an approximation of the probability of a correct majority verdict[17]:

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

where p' is the average level of competence.

In the special case in which the competence of agents is homogeneous, this is equivalent to[18]:

[15] Note that here we are still concerned with the probability of securing a majority of votes for the correct alternative – this allows us to assess the ability of a group using a judgement aggregation procedure like majority rule to identify the true state of the world. As Romeijn and Atkinson (2011) note, it is also possible to calculate the probability that an alternative is correct, given that the alternative receives a majority of the votes. Here the larger the absolute size of the majority for an alternative, the more likely that that alternative is the correct one on the agenda.

[16] Owen, G., Grofman, B. and Feld, S.L. (1989). Here N is the set of all possible combinations of votes and S is the subset of N such that the total number of correct votes is greater than ½.

[17] Grofman, B., Owen, G. and Feld, S.L. (1983) Theorem V.

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

where p is the homogeneous level of competence.

For majority voting to track the truth the competence of voters does not (necessarily) need to be *transparent* (though as we will see below, transparency helps). However if competence is not transparent then we will not know whether or not the majority verdict is reliable[19]. If the competence of agents is transparent then we can improve the epistemic performance of the aggregation procedure of majority rule by employing weighted voting. If the competencies of agents are heterogeneous, P is maximised by assigning weights to individual voters as follows[20]:

$$w_i \propto \log(\frac{p_i}{1-p_i})$$

Therefore, if we apply weights to the votes of agents according to the competence of agents, the probability of a correct majority verdict is given by[21]:

$$P = \sum_{S \subset N : \sum_{i \in S} w_i > \sum_{i \notin S} w_i} \prod_{i \in S} p_i \prod_{i \notin S} (1-p_i)$$

[18] Grofman, B., Owen, G. and Feld, S.L (1983) Theorem I.

[19] See the treatment in Dietrich, F. (2008).

[20] Grofman, B., Owen, G. and Feld, S.L. (1983) Theorem XIII.

[21] Bradley, R. and Thompson, C. (2012), adapted from Owen, G. Grofman, B. and Feld, S.L. (1989).

The epistemic success of majority voting is explained by the Condorcet Jury Theorem (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 ½ and less than 1. Formally $p = 1 > \Pr(v = x|x) > 1/2$;
- *Independence*: the probabilities 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. Formally, $\lim_{n \to \infty} P = 1$.

A simple proof for the asymptotic CJT can be found in the appendix of Dietrich (2008). It is unclear whether a proof for the non-asymptotic CJT has been published previously, but Dietrich and Spiekermann (unpublished a) includes such a proof. The classic CJT has been extended from homogeneous competence to heterogeneous competence: Theorem V of Grofman et al. (1983) states that if the distribution of individual competencies is symmetric then we obtain results analogous to the classic CJT by substituting average competence for homogeneous competence. Grofman et al. (1989) prove that in the limit, as $n \to \infty$, the asymptotic CJT holds for populations with heterogeneous competence, irrespective of the distribution of competencies, provided that the average competence is greater than 1/2. No proof of the non-asymptotic CJT for heterogeneous competencies has yet been published, and this is a weakness in the literature.

For the non-asymptotic CJT to hold for groups with heterogeneous competencies, we need a plausible interpretation of the symmetry clause in Theorem V of Grofman et al. (1983). Let p_n be the average competency of a group of n agents and P_n be the probability that a group of n agents will select the correct alternative via majority rule. The first interpretation of symmetry is that the distribution of competencies in a group with n members, $n + 2$ members, $n + 4$ members...are symmetric but that $p_n \neq p_{n+2} \neq p_{n+4}$.... In other words, the distribution of competencies is symmetric within any given group, but the average competence varies as group size varies. This cannot be the interpretation of symmetry that Grofman et al. intended as the following example shows. Suppose there is a group of three agents whose competencies are $(0.5, 0.6, 0.7)$. The average competency is $p_n = 0.6$, meaning that the competence assumption of the CJT holds, and the distribution of competencies is symmetric about the mean. The probability of this group generating a correct majority verdict is $P_n = 0.65$, so the group does better than an average member. But the individual with a competence of $p_3 = 0.7$ outperforms the group, and the non-asymptotic CJT does not hold. If we start with a group size of one agent, comprised of the agent with competence of $p_3 = 0.7$, then as we increase group size to three, the probability of a correct majority verdict is not monotonically increasing in group size.

The second interpretation of the symmetry requirement in Theorem V is that $(1/2) < ((p_1 + p_2)/2) = ((p_3 + p_4)/2) = ((p_5 + p_6)/2) = \ldots$. In other words, the distribution of competencies is symmetric within any given group, and the average competence remains constant as group size varies. But again this cannot be the interpretation of symmetry intended. Although the non-asymptotic CJT holds under this second notion of symmetry, it is an extremely restrictive condition. This notion of symmetry requires that exactly the right combination of pairs of agents is added to the group at the same time so as to maintain the average competence as group size increases. It is implausible that this would occur.

A third possible interpretation of the symmetry requirement is that agents are drawn independently from the same symmetric meta-distribution with expected value of competence >1/2. For example the meta-distribution could have a uniform distribution on $[0.2, 1.0]$. This is the interpretation that Ben-Yashar and Paroush (2000) seem to take when they modify the classic CJT. They argue that "...in reality competence is not a conspicuous characteristic of individuals and very seldom can be estimated." (p.191). Instead of the non-asymptotic CJT comparing the group epistemic performance against the competence of each individual member of the group, they compare the group epistemic performance with the expected value of an individual group member's competence. They prove that if each group member has a competence greater than 1/2 then the likelihood of a correct majority verdict is greater than the probability of a correct choice, chosen by an individual sampled at random. However this is weaker than the monotonicity of the non-asymptotic CJT which as well as implying that a group will be more reliable than an individual, also implies that a larger group will be more reliable than a smaller group.

Later in this thesis I argue that there may be cases where the competence of individual agents is *transparent*. We may know the long run accuracy of an agent's votes in a relevant reference class of social choice problems. As such we may have a good idea of the probability that they will choose the correct alternative. However the long range accuracy of an agent's votes only gives an indication as to an agent's actual level of competence. As such we cannot be certain that a given agent will in fact have the competence to outperform the group. Nevertheless there may be other cases in which agents can prove to other agents what their competence is. In these cases, the interpretations of heterogeneous competence for the CJT set out above are of no use. The non-asymptotic CJT does not hold for groups with heterogeneous and transparent competence because the probability of a correct majority verdict is not monotonically increasing in group size[22].

[22] We can see this by again considering the example of a group with competencies (0.5, 0.6, 0.7). If the group starts with the agent with $p_3 = 0.7$ then adding group members decreases the probability of a correct majority verdict.

It is important to see how the majority voting operates as a social epistemic mechanism. To do so we will need to see what types of judgement-generating factors are required for this aggregation procedure to track the truth.

Group size

There are various proofs for the classic asymptotic CJT[23]. The asymptotic CJT is often explained intuitively by the example of coin tosses[24] (and I present this explanation later in this chapter). A proof for the classic non-asymptotic CJT has only recently been presented in Dietrich, F. and Spiekermann, K. (unpublished a). Here I present an intuitive explanation of the classic non-asymptotic CJT. The following set of diagrams is intended as a pedagogical contribution to articulating how the mechanism behind the non-asymptotic CJT works.

Suppose we have an agenda with two alternatives, with one of the alternatives being correct. A vote will be taken to determine which of the two alternatives will be the social choice, and the vote will be decided by majority rule i.e. an alternative must receive more than 1/2 of the votes if it is to be the winner. We assume that the competence of voters is homogeneous and 0.6 i.e. voters have a 60% chance of voting for the correct alternative and a 40% chance of not voting for the correct alternative. The votes of any two agents are assumed to be

[23] See for example, Ladha, K. (1992) and Dietrich, F. (2008).

[24] See for example, List, C. and Goodin, R.E. (2001) and Estlund, D.M. (2008).

independent, conditional on the state of the world[25]. If we only have one agent '1′ then there are only two logically possible ways in which that agent could vote, as seen below:

Figure 2.4: the possible votes of a single agent.

	1
a	✓
b	X

'✓' means the agent voted for the correct alternative and 'X' means the agent voted incorrectly for the wrong alternative (equivalent to incorrectly not voting for the correct alternative). We can use the table above to determine the majority winner. In the first row the correct alternative received all of the votes, in the second row the correct alternative received none of the votes. The table below highlights the rows in which the correct alternative is the majority winner:

Figure 2.5: the majority winner, given a single agent.

	1	
a	✓	Winner
b	X	

Because we have an assumed level of competence, we know the likelihood of each of the logically possible outcomes, as follows:

[25] Given these assumptions, the framework of the classic CJT applies.

Figure 2.6: the probability of a majority winner, given a single agent.

	1	
a	0.6	Winner
b	0.4	

The probability that the correct alternative will be the winner is then 0.6.

But suppose we now have three agents. The logically possible combinations of votes are now as follows:

Figure 2.7: the possible combinations of votes, given three agents.

	1	***2***	***3***
a	✓	✓	✓
b	✓	✓	X
c	✓	X	✓
d	✓	X	X
e	X	✓	✓
f	X	✓	X
g	X	X	✓
h	X	X	X

If there are three voters, then for the correct alternative to be the majority winner it must receive at least two of the votes. The table below highlights the rows in which the correct alternative receives at least two votes:

Figure 2.8: the majority winner, given three agents.

	1	***2***	***3***	
a	✓	✓	✓	Winner
b	✓	✓	X	Winner
c	✓	X	✓	Winner
d	✓	X	X	
e	X	✓	✓	Winner
f	X	✓	X	
g	X	X	✓	
h	X	X	X	

As can be seen from the table there are four possible combinations of votes (i.e. ✓✓✓, ✓✓X, ✓X✓, or X✓✓) that will result in the correct alternative being the majority winner. Because we have an assumed level of competence, we can calculate the likelihood that any given combination of votes will occur. For example, the probability that all three voters vote correctly is the probability that voter 1 votes correctly AND voter 2 votes correctly AND voter 3 votes correctly. This is given by the probability that voter 1 votes correctly, TIMES the probability that voter 2 votes correctly TIMES the probability that voter 3 votes correctly. These calculations are shown in the table below, just for the rows in which the correct alternative is the majority winner.

Figure 2.9: the probability of a majority winner, given three agents.

	1	***2***	***3***	
a	0.6	0.6	0.6	$0.6 \times 0.6 \times 0.6 = 0.216$
b	0.6	0.6	0.4	$0.6 \times 0.6 \times 0.4 = 0.144$
c	0.6	0.4	0.6	$0.6 \times 0.4 \times 0.6 = 0.144$
d	0.6	0.4	0.4	
e	0.4	0.6	0.6	$0.4 \times 0.6 \times 0.6 = 0.144$
f	0.4	0.6	0.4	
g	0.4	0.4	0.6	
h	0.4	0.4	0.4	

What then is the probability that the correct alternative is the winner? There are four logically possible combinations of votes that will generate the correct alternative as the winner, so the correct alternative is the winner if it receives the votes of $1, 2$ and 3; OR it receives the votes of 1 and 2 but not 3; OR the votes of 1 and 3 but not 2; OR it receives the votes of 2 and 3 but not 1. This is given by the probability it receives the votes of $1, 2$ and 3; AND it receives the votes of 1 and 2 but not 3; and so on. This calculation is given immediately below:

$$P = 0.216 + 0.144 + 0.144 + 0.144 = 0.648$$

So with three voters the probability they will identify the correct alternative is 0.648, compared with 0.6 for an individual agent.

In fact we can see the general rule for calculating the probability that the best alternative will be the majority winner[26]. We can do this in three steps.

The probability for a given possible combination of votes occurring is given by:

$$p^h(1-p)^{n-h}$$

I.e. if there are h number of voters who vote correctly there will be $n-h$ voters who vote incorrectly. The probability of this combination of votes is given by multiplying the probabilities of the correct votes by the probabilities of the incorrect votes. But there can be several different ways of getting the same number of correct versus incorrect votes (for example to get two correct votes from three voters). To get a group of h correct votes from a wider group of agents of size n we use the following abbreviation:

$$\binom{n}{h}$$

This corresponds to:

[26] This treatment echoes Estlund (1994), but in fewer steps and reverse order.

$$\frac{n!}{h!\,(n-h)!}$$

So we have:

$$\binom{n}{h} p^h (1-p)^{n-h}$$

Finally we can specify that we want to add together all possible combinations of votes where there is a majority in favour of the best alternative. The following gives a sum of all these values:

$$\sum_{h>\frac{n}{2}}^{n}$$

So finally we have:

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

The simple calculations above illustrate the non-asymptotic version of the CJT. As the number of voters increases (from one to three) the probability that the correct alternative is the majority winner also increases. The probability that the group will select the correct alternative is monotonically increasing as the group size increases. As a group of two or more agents is greater in size than a group of one agent it follows that a group using majority rule is more likely than an individual to select the correct alternative.

To see micro-level the impact that adding more voters has on the result it is perhaps worth exploring what happens when we move from three to five voters. When we have five instead of three voters, the number of possible combinations of votes increases to 32, as shown below:

Figure 2.10: the possible combinations of votes, given five agents.

	1	***2***	***3***	***4***	***5***	***Winner***
1	0.6	0.6	0.6	0.6	0.6	Winner
2	0.6	0.6	0.6	0.6	0.4	Winner
3	0.6	0.6	0.6	0.4	0.6	Winner
4	0.6	0.6	0.6	0.4	0.4	Winner
5	0.6	0.6	0.4	0.6	0.6	Winner
6	0.6	0.6	0.4	0.6	0.4	Winner
7	0.6	0.6	0.4	0.4	0.6	Winner
8	0.6	0.6	0.4	0.4	0.4	X
9	0.6	0.4	0.6	0.6	0.6	Winner
10	0.6	0.4	0.6	0.6	0.4	Winner
11	0.6	0.4	0.6	0.4	0.6	Winner
12	0.6	0.4	0.6	0.4	0.4	X
13	0.6	0.4	0.4	0.6	0.6	Winner
14	0.6	0.4	0.4	0.6	0.4	X
15	0.6	0.4	0.4	0.4	0.6	X
16	0.6	0.4	0.4	0.4	0.4	X
17	0.4	0.6	0.6	0.6	0.6	Winner
18	0.4	0.6	0.6	0.6	0.4	Winner
19	0.4	0.6	0.6	0.4	0.6	Winner
20	0.4	0.6	0.6	0.4	0.4	X
21	0.4	0.6	0.4	0.6	0.6	Winner
22	0.4	0.6	0.4	0.6	0.4	X
23	0.4	0.6	0.4	0.4	0.6	X
24	0.4	0.6	0.4	0.4	0.4	X
25	0.4	0.4	0.6	0.6	0.6	Winner
26	0.4	0.4	0.6	0.6	0.4	X
27	0.4	0.4	0.6	0.4	0.6	X
28	0.4	0.4	0.6	0.4	0.4	X
29	0.4	0.4	0.4	0.6	0.6	X
30	0.4	0.4	0.4	0.6	0.4	X
31	0.4	0.4	0.4	0.4	0.6	X
32	0.4	0.4	0.4	0.4	0.4	X

Here, with five voters, the probability that the correct alternative wins a majority of votes[27] is $P = 0.68256$. Each of the rows in our original table in figure 2.8 has split into four sub-types. Take the first row (a) from our original table in figure 2.9. Here all voters vote correctly (and the probability for them doing so is 0.216). If we now have five instead of three voters, then there are four possibilities. The two additional voters 4 and 5 could continue the pattern and both vote correctly (figure 2.10, line 1). Or the first new voter 4 could vote correctly and the second new voter 5 could vote incorrectly (figure 2.10, line 2), or the other way round, with 4 voting incorrectly and the 5 correctly (figure 2.9, line 3). Finally it is possible that both of the two new voters 4 and 5 will vote incorrectly (figure 2.9, line 4).

The effect of adding new voters is a matter of fine graining. There is greater diversity in the logically possible combinations of votes (32 instead of 8). Just as many of the possible combinations of votes result in the correct alternative being the majority winner (exactly half). Most of the combinations of votes have no change in outcome compared with the situation with three voters. But some do - see lines 8, 12, 13, 20, 21, and 25. In line 8 for example, agents 1 and 2 vote correctly, but 3 votes incorrectly. If there were just these three voters then the correct alternative would be the majority winner (as is seen in row b of the original table). But with the addition of two voters 4 and 5, who both vote incorrectly, the result with five voters is a majority for the incorrect alternative.

[27] I present the results of the sample calculations with five decimal places so that the impact of increasing group size can be seen. Of course, it should not be thought that this level of precision is possible when applying these results to real social choice problems.

Three of the rows in the table for five voters (figure 2.10) result in a shift from a correct winner to an incorrect winner (compared with three voters, figure 2.9): lines 8 (c.f. b), 12 (c.f. c) and 20 (c.f. e). Three of the rows result in a shift from an incorrect winner to a correct winner: lines 13 (c.f. d), 21 (c.f. f) and 25 (c.f. g). In effect these 'flips' balance each other out, there are just as many 'good' flips as there are 'bad' flips. But what is interesting is that the good flips are more likely to occur than the bad flips. Compare line 8 with line 21. Line 8 delivers a bad flip (compared to line b in the table for three voters). It has two correct votes and three incorrect votes, so the probability of this combination of votes occurring is $0.6^2 \times 0.4^3 = 0.02304$. Line 21 delivers a good flip (compared with line f in the table for three voters). It has three correct votes and two incorrect votes, so the probability of this combination of votes occurring is $0.6^3 \times 0.4^2 = 0.03456$. A shift from the correct alternative losing to the correct alternative winning is more likely to occur than a shift from the correct alternative winning to the correct alternative losing. In fact if we subtract the increased probability of the correct alternative winning from the increased probability of the correct alternative losing, we get: $(3 \times 0.03456) - (3 \times 0.02304) = 0.03456$. This accounts for the increase in the probability of a correct winner between three voters and five voters i.e. $0.68256 - 0.648 = 0.03456$

Another way to see the underlying effect of the CJT is to consider the decisiveness of agents. An agent is decisive if changing her vote can change the result of the election. Let's consider the cases in which agent 1 is decisive. If there are three voters, then agent 1 is decisive in 4/8 of the cases. If there are five voters, then agent 1 is decisive in 10/32 of the cases. Increasing the number of voters decreases the importance of any given voter. Any given voter can act erratically (vote for the wrong alternative), and if there are small numbers of voters the erratic vote of a given agent will have a big influence on the final result. But if there are a large number of voters, then the erratic vote of a given voter can be weeded out by the influence of other voters. Where there are large numbers of voters, erratic behaviour (voting incorrectly) can still impact on the outcome, but for this to occur, significant numbers of voters all need to act erratically together. But where there are large numbers of voters, the probability of sufficiently large numbers of voters all acting erratically together is quite small. This is what is meant by large numbers 'weeding out noise'.

The impact on P of increasing n is illustrated in the graph below, assuming the competence and independence assumptions hold:

Figure 2.11: the probability of a correct majority verdict, as group size increases.

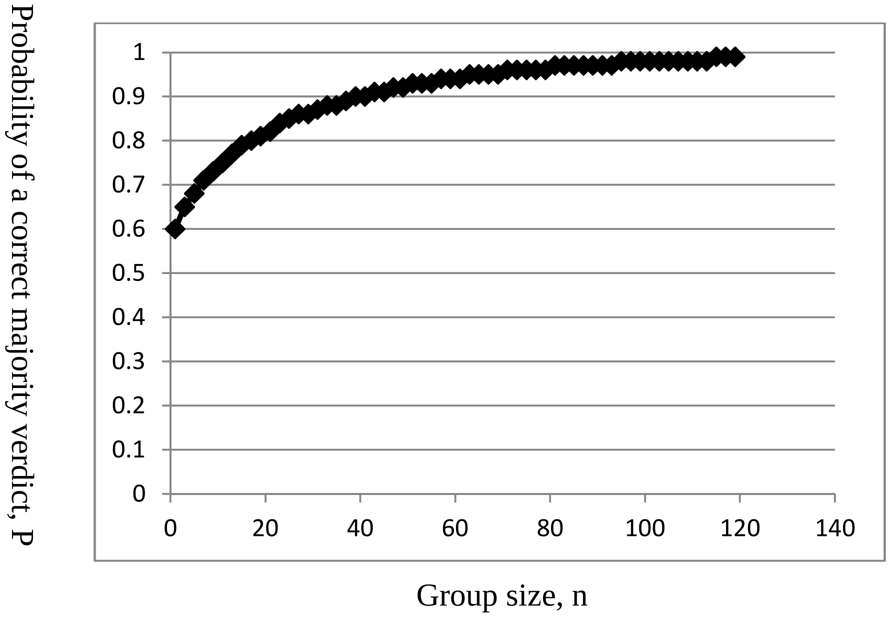

Thus far we have been addressing the non-asymptotic CJT and seen why P is monotonically increasing in n. To understand the CJT fully we need to understand why it is that for the asymptotic CJT the upper limit of P is certainty. The statistical phenomenon underlying the CJT is the law of large numbers (LLN). The LLN can be stated as follows:

> *The average value for a series of trials tends towards the expected value as the number of trials increases.*

We can see the law of large numbers in action when we toss a coin. The expected value of a fair coin is 0.5 heads. If we toss the coin a small number of times then we would not be surprised if we had something very different to 0.5 heads. But as the number of tosses increases we would increasingly expect something very close to 0.5 heads. If we had an infinite number of coin tosses we would expect exactly 0.5 of them to be heads.

The expected value of an agent voting correctly is the homogeneous level of competence p. Suppose $p = 0.6$. As group size increases it becomes increasingly likely that exactly 0.6 of the group will vote correctly. In fact if we had an infinite number of voters exactly 0.6 of them would vote correctly. 0.6 of voters in favour of the correct alternative is clearly a majority in favour of the correct alternative so as group size tends towards infinity the probability of a correct majority verdict tends towards certainty.

Majority rule, as a social epistemic mechanism, is a compensatory model of group productivity according to the Steiner (1966) taxonomy. As group size increases, the agents voting correctly offset the incorrect votes of a minority of agents.

In a political setting, majority rule may be employed for passing legislation in parliament. At least some of legislation passed in parliament can be incorrect or incorrect. For example, if the rationale for banning a certain recreational drug is that this will reduce the number of drug-related deaths then it is a matter of fact whether the change in legislation will be effective or not[28]. The United Kingdom's bicameral Westminster Parliament is comprised of the lower, democratically elected, House of Commons and the upper, appointed, House of Lords. Members of the lower house are subject to party whipping and so the

[28] *Identifying* whether the legislation has been effective or not is a different and more difficult matter.

competence and independence assumptions of the CJT do not apply and it cannot be argued that the mechanism of majority rule allows the group to track the truth. However, legislation must also be passed by the House of Lords. As of 2012, the House of Lords was comprised of 308 members of the Government and 253 members of the opposition. In addition there are 184 cross-benchers, 24 Lords Spiritual and 19 non-affiliated members. The Government's majority in the House of Lords is between 55 and 98 (depending on how the Lords Spiritual and unaffiliated members vote), If party whipping occurs and members vote according to non-epistemic grounds then, just as in the lower house, we cannot use the CJT to argue that the legislation passed by the House of Lords will be correct. However, the ostensive justification for including cross-benchers is that these members bring with them a wealth of experience from a variety of different areas of public life. The cross-bench members of the House of Lords *are* expected to cast their votes according to their best judgment. We should expect the judgements of cross-bench members to be independent (given the diverse backgrounds) and for the members to have competence levels better than random (given their experiences and successes in life)[29]. And the number of cross-benchers is sufficient to off-set the Government majority whose votes are cast on non-epistemic grounds. By employing the judgement aggregation procedure of majority rule, and by increasing the number of cross-benchers, we should expect important pieces of legislation passed by the House of Lords to be correct.

[29] The judgements of the cross-benchers may not be independent if these agents have common information. However, as Dietrich (2008) notes, independence in agent judgements can be regained by conditionalising on common factors. There may be some difficult or misleading issues that the House of Lords faces which means the average competence level of the members is less than ½. However, these issues cannot be the norm and we should expect the competence assumption of the CJT to hold in most cases.

Competence

Recall that for the CJT results to hold, the average probability of an agent correctly voting for the correct alternative must be > 0.5. Here I illustrate the impact both of competence below 0.5 and increasing competence.

In the initial calculations above, it was assumed that competence was $p = 0.6$. If individual agents instead have competence $p = 0.4$ three such agents only have a $P = 0.352$ probability of selecting the correct alternative as the social choice via majority rule. If the competence of agents is less than 0.5, then as the number of agents increases, the probability that the group will select the correct alternative via majority rule tends towards zero[30].

As the level of competence of agents increases, the probability that the group will select the correct alternative also increases. For example, three agents with a competence of $p = 0.6$ have a $P = 0.648$ probability of selecting the correct alternative. Three agents with competence of $p = 0.7$ have a $P = 0.784$ probability of selecting the correct alternative.

Independence

The independence assumption of the CJT requires that the votes of individual agents are probabilistically independent, conditional on the state of the world. Ladha (1992), Estlund (1994), and Kaniovski (2010) all consider the impact of violations of the independence assumption, of shared information and correlated

[30] See Grofman, B., Owen, G. and Feld, S.L. (1983) Theorem I.

votes. Ladha argues that the probability a majority verdict is correct is inversely related to the average correlation. 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. Kaniovski argues that a negative correlation between the votes of agents increases the probability of a correct majority verdict, while positive correlation decreases the probability of a correct majority verdict. In this section of the chapter I merely illustrate the impact that violations of independence can have on the social epistemic mechanism underlying the CJT.

There is a family of independence conditions, ranging from weaker to stronger, which capture different dependence structures. All independence conditions conditionalise on the state of the world and may then conditionalise on additional factors. The 'full blown' or standard independence assumption of the CJT requires that the votes of different agents are probabilistically independent, given the state of the world. Factors (including pieces of information) which in part determine the votes of agents, and which are held in common between agents, may lead to correlations in the votes of agents and violations of the standard independence condition. But a weaker independence condition may still be met if agents' votes are independent, conditional on any common factors and the state of the world.

We can define the following:

• $p_1 = \Pr(v_1 = x|x)$ is the probability that agent i votes for the correct alternative, given the state of the world.

• $p_{1|2} = \Pr(v_1 = x|x, v_2 = 1)$ is the probability that agent 1 votes for the correct alternative given the vote of agent 2 and the state of the world.

The standard independence assumption requires, that $p_1 = p_{1|2}$ i.e. the fact that agent 2 votes correctly or incorrectly in no way effects the vote of agent 1. This must be true for all the agents. The votes of agents will be determined by the factors they possess (including, but not limited to truth-conducive evidential and background information). If no agents have factors in common then standard independence holds. If some agents have vote-determining factors in common then standard independence will be violated. Where standard independence holds, the probability of three agents with $p = 0.6$ selecting the correct alternative via majority rule is $P = 0.648$.

Now we can see what impact violating the standard independence assumption has. Let's consider a situation in which agents 2 and 3 follow precisely what agent 1 does, because they have identical vote determining factors in common. This is illustrated in figure 2.12 below.

Figure 2.12: the probability of a correct majority winner, given dependent voters.

	1	**2**	**3**	
a	✓	✓	✓	Winner
b	✓	✓	✓	Winner
c	✓	✓	✓	Winner
d	✓	✓	✓	Winner
e	X	X	X	
f	X	X	X	
g	X	X	X	
h	X	X	X	

As we can see there are now only two possible combinations of votes: three votes for the correct alternative or three votes for the incorrect alternative. Agent 1 has a 0.6 chance of voting for the correct alternative, whereas the conditional probability of 2 voting correctly given that 1 votes for the correct alternative is 1.0. This is also true of agent 3. Therefore the probability of the correct alternative being the majority winner is:

$$P = 0.6 \times 1.0 \times 1.0 = \ 0.6$$

This result is identical to that for a single agent. The impact of violating the standard independence condition is equivalent to a reduction in the number of voters.

When agents are independent of other agents conditional on the state of the world and on common factors we can still see an effect from increasing numbers of agents. Suppose agents 2 and 3 are not independent of agent 1 given the state of the world. Formally, $\Pr(v_{2,3} = x|x) \neq \Pr(v_{2,3} = x|x, v_1)$. However agents 2 and 3 are independent of agent 1 conditional on the state of the world and common factors $C_{1,2,3}$. Formally, $\Pr\left(v_{2,3} = x\middle|x, C_{1,2,3}\right) = \Pr(v_{2,3} = x|x, C_{1,2,3}, v_1 = 1)$. The common factors lead to a bias of agents 2 and 3 that is 0.1 in the direction of agent 1's vote. So, for example, the probability of agents 2 or 3 voting correctly given that 1 has voted correctly is $+0.1$ greater than the unconditional probability[31] of agents 2 or 3 voting correctly of $p_{1,2,3} = 0.6$. Similarly the probability of agents 2 or 3 voting correctly given that 1 has voted incorrectly is -0.1 less than the unconditional probability of agents 2 or 3 voting correctly. The probability of agents 2 or 3 voting incorrectly given that 1 has voted correctly is -0.1 less than the unconditional probability of agents 2 or 3 voting incorrectly. Finally the probability of agents 2 or 3 voting incorrectly given that 1 has also voted incorrectly is $+0.1$ greater than the unconditional probability of agents 2 or 3 voting incorrectly. This is summarised in the table below.

[31] Or to be precise the probability of agents 2 or 3 voting correctly conditional just on the state of the world.

Figure 2.13: the probability of agent 2 and agent 3's votes.

		2 or 3's vote	
Agent 1's vote		✓	X
	✓	0.7	0.3
	X	0.5	0.5

This dependence transfers into the overall probability that the correct alternative will be the majority winner.

Figure 2.14: the probability of a correct majority winner, given dependencies between agents.

	1	**2**	**3**	
a	0.6	0.7	0.7	Winner $0.6\times0.7\times0.7 = 0.294$
b	0.6	0.7	0.3	Winner $0.6\times0.7\times0.3 = 0.126$
c	0.6	0.3	0.7	Winner $0.6\times0.3\times0.7 = 0.126$
d	0.6	0.3	0.3	
e	0.4	0.5	0.5	
f	0.4	0.5	0.5	Winner $0.4\times0.5\times0.5 = 0.1$
g	0.4	0.5	0.5	
h	0.4	0.5	0.5	

$$P = 0.294 + 0.126 + 0.126 + 0.1 = 0.646$$

With the standard independence assumption fulfilled, the probability of a correct winner is $P = 0.648$. If standard independence is violated and agents are not conditionally independent, the probability of a correct winner is $P = 0.6$. With violation of standard independence, but with agents independent, conditional on common factors, the probability of a correct winner drops to $P = 0.646$ in these sample calculations. This reduced epistemic performance can be interpreted as agents being less able to compensate for the mistakes of other agents, given that they share some of the vote determining factors that lead the other agents to vote for the incorrect alternatives.

It is also possible to construct examples to show how in some cases violations of standard independence are epistemically virtuous. Suppose we have a group of three agents whose competencies are as follows: $(p_1, p_2, p_3) = (0.7, 0.6, 0.6)$. As can be seen, one of the agents is more competent that the other two. Without any deference the probability of a correct majority winner is $P = 0.696$.

Now suppose that the two less competent agents defer to their more competent colleague to degree 0.1. This is represented in the figure below:

Figure 2.15: the probability of agent 2 and 3's votes

2 or 3's vote

		✓	X
Agent 1's vote	✓	0.7	0.3
	X	0.5	0.5

This dependence transfers into the overall probability that the correct alternative will be the majority winner, as follows.

Figure 2.16: the probability of a correct majority winner, given some dependence between voters.

	1	**2**	**3**	
a	0.7	0.7	0.7	Winner 0.7×0.7×0.7= 0.343
b	0.7	0.7	0.3	Winner 0.7×0.7×0.3= 0.147
c	0.7	0.3	0.7	Winner 0.7×0.3×0.7= 0.147
d	0.7	0.3	0.3	
e	0.3	0.5	0.5	
f	0.3	0.5	0.5	Winner 0.3×0.5×0.5= 0.075
g	0.3	0.5	0.5	
h	0.3	0.5	0.5	

$$P = 0.343 + 0.147 + 0.147 + 0.075 = 0.712$$

If this group of three agents vote independently the probability of a correct majority is $P = 0.696$. If they partially defer to the more competent group member the probability of a correct majority is $P = 0.712$, a clear improvement. The epistemic improvement can be interpreted as follows. Deference means agents have less ability to compensate for the errors of other agents (since agents will tend to make the same mistakes), but the deference itself generates gains in competence. The gains in competence mean agents are less likely to vote incorrectly and less likely to need other agents to compensate with correct votes. The gain in competence more than off-sets the drop in compensating ability.

It is important to note that the extent to which truth-conducive information affects competence and the extent to which shared information affects conditional probabilities of voting correctly are both big topics. More detailed treatments are given in Ladha (1992) and Estlund (1994). However, we can make four general points here. Firstly, various violations of independence mean that the probability of voting correctly is no longer equivalent to competence p_i. We can no longer use the value of an agents competence in the calculations for determining the probability of a correct majority verdict because how an agent votes depends on how other agents vote. Instead we would need to use the value of an agent's competence, given the common factors or votes of other agents (we would need to use the value of an agent's competence conditional on the common factors). Secondly, violations of standard independence are epistemically permissible (there is still value in increasing group size), provided that when we conditionalise on common factors and the state of the world there is still some randomness left in the votes of agents i.e. provided that the probability of voting correctly given the vote of another agent (or common factors) is not 0 or 1. For agents to be independent conditional on the state of the world and on common factors they need to have at least some information held uniquely by them and not shared by other agents. Being independent conditional on common factors means there is at least some chance that an agent will vote correctly when other agents vote incorrectly, so as to compensate for the incorrect votes. Thirdly, ceteris paribus violations of independence are epistemically bad. If agents are entirely dependent then as we have seen $P = p_i$ and increased group size makes no different to the probability of a correct majority verdict. Ceteris paribus, P is maximised when there are no violations of independence. Finally violations of independence can in some cases be epistemically advantageous if they increase an agent's competence level. The extent to which this is the case will depend on the assumptions made in particular cases. While violations in independence mean agent 1 is less able to compensate for the incorrect vote of agent 2, violations of independence can also increase the competence level of 2 such that 2 is less likely to vote incorrectly and less likely to need 1 to compensate with their vote.

I will consider the issue of violations of independence again in chapter 4 of this thesis, where I consider in greater detail how the judgement-generating factors of competence and independence form, and in chapter 7 of this thesis where I consider information sharing as a response to the problem of finite information.

Summary

Different aggregation procedures have the ability to track the truth *as group size increases* given certain judgement-generating factors. In other words they provide a conditional epistemic justification for group decision making. We can see a summary of the requirements in the table below.

Figure 2.17: a summary of the judgement-generating factors required for the aggregation procedures to track the truth.

	(Expert) Dictatorship	**Majority Rule**	**(Negative Reliability) Unanimity Rule**
Transparency of competence levels	Important	Not crucial, but desirable	Irrelevant
Competence levels, including distribution	Important	Important, especially distribution	Important
Independence	Harmful	Important	Important

We can make a number of comments at this point. All three aggregation procedures recommend high levels of agent competence. For us to justify the truth-tracking ability of majority rule we require a certain distribution of competencies in the group - either homogenous or symmetric about the mean - but we do not necessarily require the competencies of agents to be transparent (although this would be desirable as we could employ weighted majority rule). Dictatorship, on the other hand, does not require any particular distribution of competencies. It can be successful at tracking the truth if only one member of the group has any truth-conducive signals whatsoever. However, for dictatorship to be epistemically successful the competence of agents must necessarily be transparent. Independence is important for both majority rule (as justified by the CJT) and negative-reliability unanimity rule.

Aggregation procedures are a feature of political processes. Expert dictatorship, negative reliability unanimity rule and majority rule give conditional support to the truth-tracking ability of groups, and amount to social epistemic mechanisms. Given certain judgement-generating factors (levels, distributions and transparencies of competence; and independence of agents) groups employing these aggregation procedures can have good baseline, relative and absolute group epistemic performance. But the challenge is to account for the robustness of the antecedent of this conditional justification. We need to provide a model for how the inputs required for the epistemic success of these aggregation procedures are feasible. This is discussed in depth in subsequent chapters.

Majority voting, the CJT and law of large numbers feature prominently in epistemic defences of democracy. As such there is an extensive literature on the CJT, including extensions and critiques (some of which I covered in this chapter, some of which will be discussed in subsequent chapters). If it is to be argued that, given certain inputs, groups can use the aggregation procedure of majority voting to track the truth then we need to assess whether the literature on the CJT places further restrictions on the epistemic performance of majority voting as an aggregation procedure. We begin this task in the next chapter. Chapter 3 focuses on the challenge that increased agenda size poses for the CJT.

Chapter 3: The Condorcet Jury Theorem - agenda size and competence.

The discussion in the previous chapter on judgement aggregation procedures assumed that the agenda for a social choice comprised two alternatives, one correct and one incorrect. An obvious concern with this simplification is that very many real-world social choice problems have more than two plausible alternatives, and a social planner may not be able to identify the one correct alternative in advance in order to set a two-placed agenda. In addition, increasing agenda size may decrease the competence level of agents, the probability that individual agents will vote for the correct alternative.

Increasing agenda size poses particular problems for the judgement aggregation procedure of majority rule and the classic Condorcet Jury Theorem. In this chapter I consider extensions of the classic Condorcet Jury Theorem to cope with multiple alternatives. I argue that existing extensions to the classic Condorcet Jury Theorem run the risk of either requiring too much effort on the part of the agents, or they risk decreasing the competence of agents. I argue for a mixed approach for extending agenda size beyond two alternatives, with multiple elections, agendas containing multiple alternatives and varying group sizes.

I also address the challenge posed to the Condorcet Jury Theorem by David Estlund's Disjunction Problem. The Disjunction Problem makes use of varying agenda size to contest whether it is possible, in principle, to justify the competence assumption holding. I carefully analyse the Disjunction Problem and argue that at best the Disjunction Problem shows how the framing of an agenda by a social planner can impact on whether the competence assumption does or does not hold. However the Disjunction Problem, as a criticism of the Condorcet Jury Theorem, relies on a straw-man argument. No one should attempt such an in principle justification for the competence assumption holding in particular social choice problems.

Agenda size

Very many social choice problems will have more than two possible alternatives. There can also be an objective quality ordering over that set of alternatives: some of the alternatives will be good, some will be bad and some will be neutral. For example if a group has a unique agreed goal of reducing carbon emissions, then a poor policy alternative would be to reduce the tax on petrol (people will drive more if driving is cheaper, increasing emissions). A policy alternative of encouraging hydrogen powered cars may be neutral if the hydrogen fuel is produced from fossil fuels. A policy alternative of encouraging fuel efficient cars may be a good alternative, but not as good at reducing carbon emissions as an alternative of producing all energy via nuclear fission (see figure 3.1 below).

Figure 3.1: a quality ranking of alternatives.

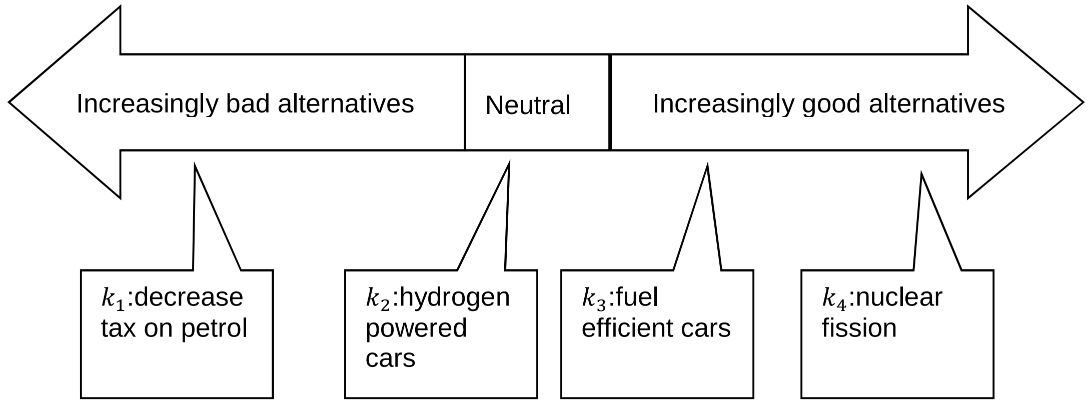

The set of possible alternatives and the objective quality ordering over that set are matters of fact. But the issue we are dealing with is an epistemic one. Agents do not have immediate access to this objective quality ordering and this can be for a variety of reasons. Firstly, the alternatives may not be presented in a ready-made quality ordering and alternatives could be mixed up with non-alternatives. Agents need to do some work to separate out the relevant alternatives (for example those that could have an effect on carbon emissions) from those alternatives that are irrelevant to the issue at hand (for example, alternatives that are more to do with improving health care). Once agents have identified and separated out the alternatives they also need to sort them into their appropriate quality ordering if they are to select the best one[32]. But this initial level of opaqueness may be relatively easy for individual agents to deal

[32] To be clear the concern of this thesis is social choice problems where there is one uniquely correct alternative. As such, even when there is quality ordering over a set of alternatives, it is not necessary for a group to identify the complete quality ordering. It is only necessary to identify which alternative is strictly better than all the others.

with, or it may be something that a social planner (such as an electoral official) can do in setting an agenda.

The second level of epistemic difficulty arises because the intrinsic quality of alternatives may not be clear and the sheer number of alternatives may create further opaqueness. Even experts may have some difficulty in determining the intrinsic quality of some alternatives. For example, the extent to which fuel efficient cars will reduce emissions depends on a variety of factors including the design of the car engines, the price of fuel and how people respond to incentives. As a result it can be very difficult to predict the outcome of such a policy. Also, it seems plausible that a given individual's ability to identify the best alternative decreases as the number of alternatives increases. Whereas an individual might be quite competent at identifying the best alternative from a set of five, they may have more difficulty at selecting the best alternative out of a set of 100, 1000 or 100,000. Arguably the following thesis is prima facia plausible:

> *As the number of alternatives tends towards infinity the competence of agents, the probability that they will vote for the best alternative, tends towards zero.*

The extent to which this thesis is true - the degree to which competence decreases as the number of alternatives increases - is an empirical matter and is likely to depend on the contingent circumstances of a particular social choice problem. But consider the example from figure 3.1 above regarding alternatives for dealing with climate change. Agents might have a competence of $p = 0.3$, they may have a 30% chance of voting for the correct alternative of nuclear fission. This uniquely correct alternative could be further refined into more specific alternatives. Let's call these finer grained alternatives for nuclear fission ‘reactor type A’, ‘reactor type B’, ‘reactor type C’ and ‘reactor type D’. The objective quality ordering over these alternatives is: reactor type A > reactor type B > reactor type C > reactor type D. The probability of voting for one of the finer grained alternatives of nuclear fission must sum to 0.3; the 0.3 probability of voting for the correct alternative must be divided among the finer grained alternative when they are placed on the agenda. So unless agents have zero probability of voting for reactor type B, reactor type C and reactor type D, then as the correct alternative is refined into more specific alternatives, competence drops. This second level of epistemic opaqueness provides a potential niche for employing the epistemic power of groups.

With the aggregation procedure of majority rule the problem of increasing agenda size is particularly acute. If agenda size increases, the competence of agents can decrease. If the competence of agents drops below a half (i.e. $p < 0.5$) then the probability of an agent voting for any one of the *incorrect* alternatives will be greater that 0.5. The classic CJT states that when $p < 0.5$ the probability of a correct majority verdict is monotonically decreasing in group size and in the limit tends to zero[33]. In other words, the increasing number of alternatives and associated impact on agent competence levels means

[33] See Grofman, B., Owen, G. and Feld, S.L. (1983) Theorem I.

majority voting harms a group's ability to track the truth; increasing group size is an epistemic disadvantage.

The problem is that the classic CJT is simply silent on social choice problems where there are more than two alternatives. The set of possible political social choice problems that only involve two alternatives is likely to be a very small subset of the set of all possible political social choice problems. Therefore, thus far, we only have a very limited justification for majority voting as a mechanism by which groups of agents can track the truth.

Here I outline a number of extensions of majority rule and the classic CJT to cope with more than two alternatives. Firstly, there are the Condorcet and Borda extensions, as discussed by Young (1988). The other major extension comes from List and Goodin (2001) who extend the classic CJT to plurality rule over many alternatives. I will discuss each of these in turn.

<u>Condorcet and Borda extensions of the dichotomous CJT</u>

Here I follow the treatment in Young (1988), where the author goes to considerable effort to clarify and reconstruct Condorcet's own approach for extending majority rule beyond two alternatives.

Condorcet's own extension of majority rule to cover multiple alternatives requires that each possible combination of alternatives is voted on in a pairwise fashion[34]. For example, if there are three alternatives k_1, k_2 and k_3 then three elections need to occur: $k_1\, vs.\, k_2$, $k_2\, vs.\, k_3$ and $k_1\, vs.\, k_3$. If an alternative (k_1, k_2 or k_3) is the majority winner in every pairwise comparison then it is likely to be the best alternative.

It is unclear why pairwise sequential voting should track the truth. Young (1988) is primarily concerned with procedures for identifying the correct quality ordering or ranking rather than identifying the one best or correct alternative. However here I provide a brief justification for why Condorcet's own extension to multiple alternatives is likely to generate the correct winner in a series of pairwise choices. Suppose that the objective quality ordering over alternatives is $k_1 > k_2 > k_3$, meaning that k_1 is the uniquely best alternative. The probability that k_1 will be the majority winner ($\mathrm{Pr}_M(k_1)$) is given by:

$$\mathrm{Pr}_M(k_1) = \mathrm{Pr}_M(k_1 > k_2) \times \mathrm{Pr}_M(k_1 > k_3)$$

In other words the probability that alternative k_1 will be the majority winner in all pair-wise comparisons is given by the probability that k_1 will be the majority winner in a comparison against k_2 ($\mathrm{Pr}_M(k_1 > k_2)$) times the probability that k_1 will be the majority winner in a comparison against k_3 ($\mathrm{Pr}_M(k_1 > k_3)$).

Similarly the probabilities that the incorrect alternatives k_2 or k_3 are the majority winners over all other alternatives are as follows:

[34] Note again that we are not interested in generating a complete quality ordering over all the alternatives, just identifying the uniquely best alternative. Therefore we ignore Condorcet's and Borda's proposals for generating a complete quality ordering over multiple alternatives.

$$\Pr_M(k_2) = \Pr_M(k_2 > k_1) \times \Pr_M(k_2 > k_3)$$
$$\Pr_M(k_3) = \Pr_M(k_3 > k_2) \times \Pr_M(k_3 > k_1)$$

If the competence and independence assumptions of the classic CJT hold then it is more likely that there will be a correct majority winner than an incorrect majority winner in each pair-wise choice. Where the two CJT assumptions hold, the probability of a correct majority winner will be $P > 0.5$. Where the competence and independence assumptions hold, the probability of an *incorrect* majority winner is $P < 0.5$. Therefore the probability of the three different alternatives being the majority winner in every pairwise comparison is respectively:

$$\Pr_M(k_1) = (> 0.5) \times (> 0.5)$$
$$\Pr_M(k_2) = (< 0.5) \times (> 0.5)$$
$$\Pr_M(k_3) = (< 0.5) \times (< 0.5)$$

As can be seen the correct alternative k_1 is more likely to win a majority in all pairwise comparisons than any of the other alternatives are. It is unclear what to do when no alternative wins every pairwise comparison or where there are cyclical majorities. Young proposes a modified solution to Condorcet's own proposal for resolving cyclical majorities when we are interested in *complete rankings*. Young's solution is to reverse the ordering of the pairwise comparison of alternatives that have the least combined plurality. If we apply this solution to the problem of finding the *uniquely best alternative* then to overcome cyclical majorities we pick the pairwise comparison that had the lowest margin of victory and reverse the verdict. This is justified by Grofman et al. (1983) Theorem III "For $p > 0.5$, the larger the size of the majority in favour of an alternative, the more likely is that alternative to be the correct one" (p.265). If no alternative wins every pairwise comparison we should reverse the result with the lowest margin of victory, since this is least likely to be correct.

We can use Condorcet's own approach to ensure that, where there are multiple alternatives, every possible alternative is considered by the voting group. As a consequence the resultant social choice is almost certainly likely to be the best alternative (assuming the competence and independence assumptions of the CJT are met and that there are sufficient numbers of voters). It follows that majority rule continues to meet the standards of baseline, relative and absolute epistemic performance for social choice problems involving multiple alternatives. As group size tends towards infinity the uniquely best alternative is almost certain to be selected by the group and the group is more likely than an individual or a random choice to select the best alternative.

But there is a problem with this extension of the CJT to underwrite the epistemic performance of majority rule. It demands that we have a pair-wise comparison between every possible combination of alternatives. With even modest numbers of alternatives we will require a lot of elections. The precise formula for determining the number of elections is:

$$e = \frac{k(1 - k)}{2}$$

where e is the number of elections and k is the number of alternatives. So if, for example, there are 100 possible alternatives we will need 4,950 elections. If there are 1000 possible alternatives then we will need 499,500 elections. Even nearly five thousand elections (for one hundred possible alternatives) is a significant burden to place on voters.

A simpler way of achieving a pair-wise comparison between every alternative is not to run $e = \frac{k(1-k)}{2}$ elections but rather to ask agents to provide their judgement of the quality rankings over the alternatives. For example, an agent could rank the alternatives (correctly) as $k_1 > k_2 > k_3$. From this information it is clear that if this agent were forced to make a series of pairwise comparisons between each of the alternatives they would cast a vote for $k_1 > k_2$, $k_1 > k_3$ and $k_2 > k_3$. The advantage of asking for rankings rather than a series of pairwise judgements is that it requires less effort on behalf of agents: agents only need to fill out one (possibly lengthy) ballot sheet. There is little point in agents participating in a pairwise comparison of $k_1\ vs.k_3$ when it has already been determined that they would vote for $k_1 > k_2$ and $k_2 > k_3$. Moreover, asking each agent for a ranking of alternatives avoids possible inconsistency on behalf of the agent, for example it prevents an agent from voting for $k_3 > k_1$ when they have already voted for $k_1 > k_2$ and $k_2 > k_3$.

If agents provide a social planner with a complete quality ordering then it is possible to use an aggregation procedure other than the Condorcet winner. Instead we can employ a Borda count. Under a Borda count, the alternative that is ranked last in an agent's ordering receives no points. The alternative that is ranked second to bottom receives one point, the alternative ranked third to bottom receives three points, and so on. The winning alternative is the one that receives the most points. Importantly, the Borda winner is more likely to be the correct social choice than the pair-wise Condorcet winner is[35].

Even if we choose to employ a Borda count to cope with multiple alternatives, this can still place significant burdens on the voters. It takes little time to rank three alternatives in order of quality. Sorting 25 possible alternatives into the apparent quality ordering may take considerable time and effort. Ranking 100 or 1000 alternatives would be an unfair task to impose on most agents.

The beauty of the standard CJT framework of a dichotomous choice is that agents are only required to cast one judgement, namely for the alternative they judge to be the best. Fortunately there is an extension to majority voting to cope with multiple alternatives that preserves this simplicity.

<u>The List and Goodin extension of the CJT to plurality voting over multiple alternatives</u>

List and Goodin (2001) extend the classic CJT from majority voting over two alternatives to plurality voting over many alternatives. Under the aggregation procedure of plurality rule, an alternative is the social choice if and only if it receives strictly more votes than any of the other alternatives. As discussed in

[35] See Young (1988). According to Young's arguments if agent competence is high and group size is sufficiently large then any reasonable decision rule is likely to generate the correct social choice. However if competence is close to 0.5, the Borda count is epistemically superior to a pairwise Condorcet choice.

the previous chapter, the classic CJT result relies on the law of large numbers. Suppose we have a fair coin. In the long run we should expect this coin to give us heads roughly half of the time. If we were to toss the coin just a few times - say ten times - then we would not think it unusual to get 6, 7 or 8 heads. But if we toss the coin a thousand times then we are very unlikely to get 600, 700 or 800 heads. As the number of tosses increases it becomes increasingly likely that we will get heads half of the time. The competence assumption of the CJT requires that on average voters select the better of two alternatives slightly more than half of the time. Therefore as the number of voters increases it becomes increasingly likely that the better of two alternatives will receive slightly more than half of all votes. If the better of two alternatives receives slightly more than half of the votes it will be the majority winner. Therefore as the number of voters increases it becomes increasingly likely that the better of two alternatives will be the majority winner.

The List and Goodin result also rests on the law of large numbers and flows naturally from the classic two-alternative case. Suppose we now have three alternatives - k_1, k_2 and k_3. k_1 is the best alternative and voters have a competence of $p = 0.4$ i.e. agents have a 40% probability of correctly voting for k_1 as the best alternative. Voters have a 0.3 probability of selecting k_2 and a 0.3 probability of selecting k_3 as the best alternative. As the number of voters increases it becomes increasingly likely that k_1 will receive 40% of the vote whereas k_2 and k_3 will receive 30% of the votes each. k_1 will not be the majority winner because k_1 receives 40%, less than the >50% required to be the majority winner. $\text{Not} - k_1$ $(\neg k_1 \equiv (k_2 \vee k_3))$ receives 60% of the vote and so if this was a simple pair-wise comparison between the correct k_1 and incorrect $\neg k_1$, the incorrect $\neg k_1$ would be the majority winner. But k_1 will be the plurality winner as k_1 receives more votes than any other alternative (on the three-placed agenda k_1 receives 40% of the votes, which is more than the 30% k_2 receives and more than the 30% k_3 receives).

We can use the List and Goodin extension of the CJT to ensure that every possible alternative is considered by the voting group, by placing every possible alternative on the same agenda. We can state the extended CJT as follows. Suppose we have a social choice problem where plurality voting will be used to identify the correct alternative as the social choice from a set of k possible alternatives. The extended CJT has two assumptions:

- *Extended competence*: the probabilities that agents will vote for the correct alternative are homogeneous and greater than the probability that they will vote for any of the other alternatives;
- *Independence:* the events of any two agents voting for the correct alternative are independent.

The extended CJT[36] result comes in two parts:

- *Non-asymptotic extended CJT*: the probability that the group will select the correct alternative is monotonically increasing as the group size increases;

[36] List and Goodin (2001) focus on the extended asymptotic CJT and provide a formal proof for this result. The extended non-asymptotic CJT is illustrated with sample calculations in their paper.

• *Asymptotic extended CJT*: in the limit as group size tends towards infinity, the probability of a correct plurality verdict tends towards certainty.

Plurality rule then becomes a replacement for majority rule as a democratic aggregation procedure[37]. Plurality rule exhibits the virtues of good baseline, relative and absolute epistemic reliability. Provided that the extended competence and independence assumptions hold and provided the group is of a sufficient size then a group employing plurality rule is more likely than random, more likely than an individual and likely simpliciter to select the correct alternative as the social choice[38].

The use of the List and Goodin extension of the CJT to support the epistemic performance of democratic aggregation procedures runs up against the problem discussed earlier in the chapter, namely as the number of alternatives increases the competence of agents may decrease. Note that, unlike in the classic CJT, the problem is not that competence levels will drop to a level such that the extended CJT competence assumption does not hold. Rather the concern is that because competence levels are so low, the probability of a correct plurality winner will be too low.

[37] Plurality rule is very similar to majority rule. In both cases every agent gets to cast a vote. In both cases, the vote of more than one agent but fewer than every agent is required to determine the social choice. The difference between plurality and majority voting is over the threshold for an alternative being the social choice. Majority voting requires strictly more than half of all votes be in an alternative's favour if it is to be the social choice. Plurality voting has a lower threshold for determining the social choice, it merely requires than an alternative receive more votes than any other alternative if it is to be the social choice.

[38] List and Goodin (2006) generalise May's Theorem from majority rule to plurality rule. May's Theorem shows that majority rule is the only aggregation procedure that satisfies four important democratic virtues including universal domain, anonymity, neutrality and positive responsiveness. With the two List and Goodin results (2001, 2006) we have good reason for claiming that plurality rule should be the default democratic aggregation procedure: it can track the truth and it preserves important democratic virtues.

There are many ways in which competence could decrease as the agenda size increases. The rate at which competence decreases is likely to depend on the type of social choice problem. Here I will consider one possibility for how competence decreases as agenda size increases. Under the 'ratio' rule, the ratio of competence to the probability of voting for a given incorrect alternative remains constant irrespective of agenda size[39]. The basic intuitive justification for the ratio rule is that an agent's ability to identify the best alternative (their competence) depends on their ability to distinguish correct from incorrect alternatives. Agents' probabilities of voting for the correct alternative remain proportional to their probabilities of voting for the incorrect alternatives.

We will assume that the competence assumption of the extended CJT holds, so that probability of agents voting for the correct alternative is greater than the probability of them voting for any of the incorrect alternatives. Furthermore we can make the simplifying assumption that the probabilities of voting for any of the incorrect alternatives are identical (that error is distributed evenly over the incorrect alternatives). This is stricter than the List and Goodin competence assumption, which only requires that the probability of voting for the correct alternative is greater than the probability of voting for any of the wrong alternatives, and which allows for the probabilities of voting for the different wrong alternatives to vary. This simplifying assumption is made both for the sake of rhetorical clarity and for the sake of ease of calculations.

[39] I also assume that error is distributed evenly over the incorrect alternatives so that the probability of voting for one incorrect alternative is identical to the probability of voting for any other given incorrect alternative.

If p_1 is the competence of agents (the probability that they will vote for the correct alternatives) and p_2 is the 'incompetence' of agents (the probability that they will vote for each of the incorrect alternatives) then the following ratio remains constant irrespective of agenda size[40]:

$$\varepsilon = \frac{p_1}{p_2}$$

Because of our assumption that the probabilities of voting for any of the incorrect alternatives are identical, we know that:

$$p_2 = \frac{1 - p_1}{k - 1}$$

where k is the number of alternatives. Therefore[41]:

$$\varepsilon = \frac{(k-1)p_1}{1 - p_1}$$

There could be infinitely many different variations of the 'ratio' rule because there are infinitely many possible values for the ratio ε. Again, the extent to which voter competence decreases as the number of alternatives increases is an empirical matter and will depend on the contingent circumstances of particular social choice problems. Nevertheless, if it is the case that voter competence

[40] Note that $p_2 \neq 1 - p_1$ but rather $p_1 + (k-1)p_2 = 1$.

[41] Equivalently, $\frac{p_1}{1-p_1} = \frac{\varepsilon}{1-k}$. Here we can see that agent competence levels on a given agenda depend both on the number of alternatives on the agenda *and* on the value of the ratio variable ε.

does decrease as the number of alternatives increases, then this flows through into the probability that the best alternative will be selected as the social choice. If the number of alternatives k increases, then the competence of agents decreases. If at the same time the number of agents n remains constant then the probability P of a correct plurality verdict decreases. We can counteract the drop in the probability of a correct plurality verdict, caused by increasing numbers of alternatives and the associated drop in agent competence levels, by increasing the number of agents. The extended CJT implies that as we increase the number of agents (who have a fixed level of competence, given the number of alternatives on the agenda) then P increases. But there will come a point at which the numbers of agents required to compensate for the drop in competence levels exceeds the number of voters that can reasonably be assumed to exist.

A mixed approach to extend the CJT

In sum, the problem with applying the earlier extensions to majority rule of Condorcet and Borda is that they place too much of a burden on voters; the problem with applying the List and Goodin extension is that it requires more voters than can reasonably be assumed to exist. As such these extensions of the CJT, which allow for more than two alternatives on the agenda, each face practical problems. However, we can use the insight gained from these extensions to show that it is feasible to increase the number of alternatives on an agenda and still retain the truth-tracking ability of democratic aggregation procedures. The Condorcetian extension (as presented by Young) adjusts the variable of the number of elections. The List and Goodin extension adjusts the variable of the size of the agenda. Although not canvassed in either extension, we could also adjust the number of voters involved with each election. In effect, we have three variables at our disposal which we can use to balance out the challenges posed by increased numbers of elections, increased agenda size and finite voters.

If a series of pairwise elections will require too many elections, we can increase the size of each agenda to reduce the number of elections. We can also reduce the burden on individual agents by reducing the group size in each election so that an agent only has to cast a vote on some (not all) of the agendas

If a large agenda size means that competence will drop too low, we can boost competence by reducing the size of the agenda and having more elections.

As I have stressed previously, the extent to which voter competence decreases as the number of voters increases is an empirical matter. It is also an empirical matter as to how many alternatives there are and how many agents there are. The trade-offs between the number of voters, number of elections and size of agendas will therefore need to be made on a case by case (or type by type basis) and more general results will be questionable. However I will now provide a set of examples and sample calculations to show how this mixed approach to extending the CJT to multiple alternatives can be superior in some circumstances to both the Young and the List and Goodin extensions.

A comparison of CJT extensions via sample calculations

Suppose we have 25 possible alternatives, one of which is objectively the best and 1001 voters. We have four approaches for generating the social choice, and we are looking for the approach that is most likely to select the objectively best alternative as the social choice (while minimising the burden placed on voters). The four approaches are: majority voting and the classic CJT; the Condorcetian extension; the List and Goodin extension; and my mixed approach.

Majority voting and the classic CJT

Here we are restricted to a simple pairwise vote between two alternatives. Because there are only two alternatives I will make the simplifying assumption that voter competence $p = 0.6$. The probability that the best alternative (from a set of 25 alternatives) will be the social choice is the probability that the best alternative makes it on to the agenda, multiplied by the probability that the best alternative will be selected in the pairwise vote (given that it is on the agenda). If there is no special way of setting the agenda, then we can assume that it is a random chance that a given alternative will secure a place on the agenda[42]. The probability that any given alternative (including the objectively best alternative) makes it onto the agenda is $((1/25) + (1/24))$. The probability that 1001 voters with $p = 0.6$ select the better of two alternatives on an agenda is[43] $P = 1.0$.

[42] Of course it is implausible that the two places on the agenda would be set by a random lottery. Nevertheless there remains a question of how the agenda for a single pairwise choice would be set, how multiple alternatives can be weeded down to the two alternatives permitted by the classic CJT. I address the problem of agenda setting later in the thesis.

[43] The value of 1.0 is approximate and involves rounding. Under the standard CJT the probability of a correct majority verdict approaches (but does not reach) certainty.

Therefore, if we employ standard majority rule and the classic CJT, then the probability that the best alternative will be the social choice is $((1/25) + (1/24)) \times 1.0 = 0.082$. The group has roughly an 8% chance of selecting the best alternative in this example where there are multiple alternatives. As such, where there are multiple alternatives, majority voting is slightly better than random at selecting the correct alternative (8% rather than 4%); a group employing this aggregation procedure is slightly better than an individual at selecting the best alternative (8% rather than 4.9%); but an 8% probability of selecting the correct alternative is surely far too low for the group to be considered likely to identify the correct alternative. As such majority voting lacks absolute epistemic performance.

The Condorcetian extension

The Condorcetian extension, on a strict interpretation, requires that we consider each possible combination of alternatives in a pair-wise fashion. If we do this then the number of elections required is:

$$\left(\frac{25(25-1)}{2}\right) = 300$$

Arguably 300 elections are too much of a burden to place on each of our 1001 voters. We could reduce the burden placed on voters by instead asking them to provide a complete ranking over the 25 alternatives and then either looking for the alternative that is a majority winner in every pair-wise choice (which still requires the social planner to consider the outcome of 300 elections), or the social choice could be determined by a Borda count. Whichever aggregation procedure is employed, requiring agents to provide a complete quality ordering over 25 alternatives is still a significant burden, given the minimal requirements placed on voters by plurality rule.

If we employ the Condorcetian pairwise criterion, then the probability of a correct social choice is $P = 1.0$.

The List and Goodin extension

Here we have one election with 25 alternatives where we look for a plurality winner. Because we now have 25 alternatives on the agenda rather than 2, I am assuming that voter competence at identifying the best alternative will decrease from the 0.6 value seen in the classic CJT case, in line with the ratio rule. For the majority voting case with two alternatives on the agenda $p_1 = 0.6$, $p_2 = 0.4$. Therefore:

$$\varepsilon = \frac{p_1}{p_2} = \frac{0.6}{0.4} = 1.5$$

The competence for the 25 alternative case is therefore[44]:

$$\varepsilon = \frac{p_1}{p_2} = \frac{0.05882}{0.03992} = 1.5$$

Given these assumptions, where the probability of an agent voting for the correct alternative is $p_1 = 0.05882$ and the probability of any agent voting for

[44] Notice that $p_1 + 24(p_2) = 0.05882 + 24(0.03992) = 1$.

each of the incorrect alternatives is $p_2 = 0.03992$, the probability that a group of 1001 voters will select the best alternative using plurality rule is 0.766. A group using plurality rule is more likely than random (4%) and more likely than an individual agent (6%) to select the correct alternative. But the 77% probability of selecting the correct alternative is still some way from a group being guaranteed to select the correct alternative as the social choice.

A mixed approach

Here we can employ the insights gained from the Condorcetian and List and Goodin extensions. We have at our disposal many possible combinations of numbers of elections, sizes of agendas, and sizes of voters cohorts which we can use to simultaneously balance the burden placed on voters and the need to increase the probability that the best alternative will be the social choice. One approach is to split the 25 alternatives into 5 groups of 5 alternatives and hold plurality elections on each of these sub-agendas involving all voters. The winning alternative from each of these sub-elections can then go forward for a final agenda to select the social choice. This is illustrated in figure 3.2 below.

Figure 3.2: a mixed approach for multiple alternatives.

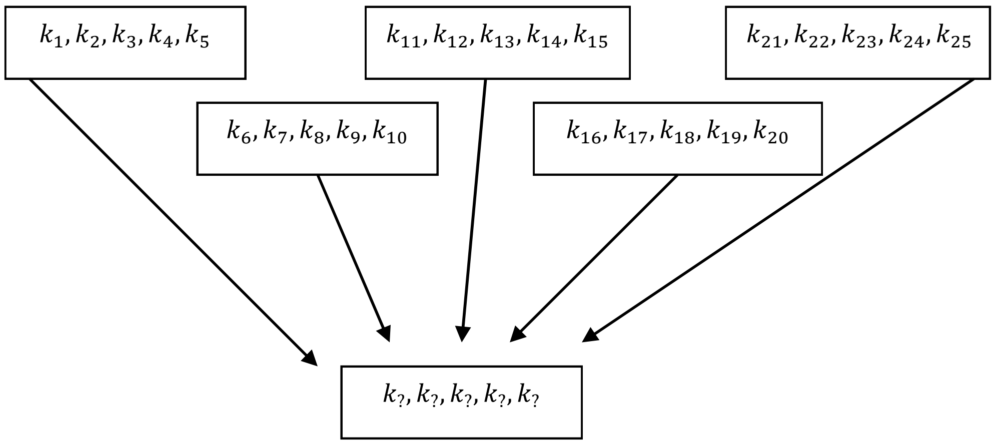

Because we have agendas comprising 5 alternatives, competence according to the ratio rule is as follows[45]:

$$\varepsilon = \frac{p_1}{p_2} = \frac{0.27272}{0.18181} = 1.5$$

Given these assumptions, the probability that a group of 1001 voters will select the best alternative from an agenda of 5 alternatives is[46] $P = 1.0$. And the probability that the best alternative, from a set of 25 alternatives, will be the social choice is given by the probability that it is selected in the first sub-election, multiplied by the probability that it is selected in the final election: $1.0 \times 1.0 = 1.0$. With the mixed approach employed here we are close to

[45] And notice also that $p_1 + 4(p_2) = 0.27272 + 4(0.18181) = 1.0$.

[46] Here again this value of 1.0 is approximate and involves some rounding.

certain to select the correct alternative as the social choice and agents only have to cast votes in 6 elections to generate it. This mixed approach has demonstrated baseline, relative and absolute epistemic performance: the mixed approach is more likely than random, more likely than an individual and likely simpliciter to generate the correct alternative as the social choice.

In sum, we have at least four ways of coping with an agenda greater than two alternatives. The performance of the different approaches for dealing with multiple alternatives is summarised below. I include both the probability of a correct social choice (P) and the number of elections (e) required for each approach.

Figure 3.3: a summary of the approaches for coping with multiple alternatives.

	P	***e***
Majority voting	0.082	1
Condorcetian extension	1.0	300
List and Goodin extension	0.766	1
A mixed approach	1.0	6

A mixed approach, of varied numbers of elections, agenda size and group size is the most successful at preserving the epistemic performance of democratic aggregation procedures like majority or plurality voting (while at the same time reducing the burden on voters). The epistemic performance of the mixed approach could be improved even further. If the competence of agents is heterogeneous and transparent we can task different agents with voting on the agenda that they are most competent on, and prevent them from voting on agendas where their competence is low or they have a bias. We may also be able to set the sub-agendas in such a way as to make the quality of alternatives more transparent (for example by avoiding placing similar alternatives on the same agenda).

It is important to stress that the precise combinations of numbers of elections, agenda sizes and voter cohort sizes that are optimal is a contingent matter. Given the assumptions of initial voter competence of 0.6 and the ratio rule, we get the results in figure 3.3 that advocate a mixture of six elections on agendas of five alternatives using plurality rule. However, if the competence of agents is more resilient to increases in agenda size then a single election on an agenda containing all the possible alternatives (the List and Goodin 2001 model) will do just as well as a mixed approach.

To know what combination of agenda size and number of elections is optimal requires knowing something about how competence degrades as the number of alternatives increases. Having multiple elections generates little epistemic loss but does place additional burdens on agents. When there are multiple elections, as per the mixed approach, then the probability that the correct alternative will be the social choice is the probability that the correct alternative wins each election where it is placed on the agenda. Therefore the probability of a correct social choice is the *product* of the probabilities that the correct alternative wins each election. As the number of elections facing the correct alternative increases, the probability that it will be the final social choice *decreases*. However, provided that the probability the correct alternative will win each election is high (because of high agent competence levels on a small agenda), the product of these probabilities is high and extra elections do not pose a significant epistemic problem. There are however clear epistemic gains to multiple elections with smaller agendas. By allowing the same agents to vote on several agendas we in effect increase the number of high competence agents in the group.

The disjunction problem

The standard extension to the classic CJT to cope with agendas of more than two alternatives is the List and Goodin extension. This extension addresses the obvious criticism of the classic dichotomous CJT, that it is only applicable in limited circumstances. But the extension of the CJT to agendas of multiple alternatives may generate further vulnerability. The 'Disjunction Problem' makes use of the extension of the CJT to multiple alternatives to challenge the fulfilment of the competence assumption[47].

[47] For clarification, the discussion in the first half of this chapter considered the problem of how competence levels can decrease as agenda size increases. It was assumed that the competence assumption of the extended CJT still held no matter how many alternatives there were (that the probability of voting for the correct alternative was still greater than the probability of voting for any other alternative, no matter how many alternatives there were). The Disjunction Problem, by contrast, argues that we have no justification for the extended CJT competence assumption holding.

David Estlund sets out the Disjunction Problem in chapter 12 of his 2008 *Democratic Authority: a Philosophical Framework*. The book as a whole argues in favour of an epistemic proceduralist defence of democracy. According to epistemic proceduralism, political decisions are legitimate and agents are obliged to follow them because the procedures that generated these decisions tend to produce correct decisions. For the thesis of epistemic proceduralism to hold, it is necessary to provide a mechanism or justification for why it is that political or democratic decisions tend to be correct. Majority or plurality voting and the CJT would seem like a natural fit with epistemic proceduralism, but Estlund provides a series of criticisms of the CJT as a mechanism to underwrite epistemic proceduralism. The Disjunction Problem is a new criticism. My concern here is not epistemic proceduralism, but is rather whether the Disjunction Problem really is a problem for the CJT.

To present the Disjunction Problem clearly it is helpful to fill in some of the detail missing in the exposition provided by Estlund. Firstly Estlund argues that the CJT competence assumption requires that agents have a level of competence that is better than random. Suppose we have k alternatives on an agenda. Agents' homogeneous level of competence, the probability that they vote for the correct alternative, is represented as p. Random competence, the probability that an agent would vote for the correct alternative on an agenda by chance, is defined relative to the agenda size. So if there are k alternatives on the agenda random competence is $p = 1/k$. Firstly let's consider the classic dichotomous CJT. For the classic dichotomous CJT $k = 2$ and $p > 1/k$ if and only if $p > 1/2$. Therefore the classic CJT competence assumption does indeed require better than random levels of competence. However things are different for the extended CJT. For the extended CJT, $k \geq 2$ and the competence assumption requires that the probability of voting for the correct alternative is greater than the probability of voting for any of the other alternatives. If the competence levels of agents are less than or equal to random (if $p = 1/k$ or $p < 1/k$) then it is *not* the case that agents are more likely to vote for the correct alternative than for any of the incorrect alternatives and the competence assumption does not hold. Therefore if the extended competence assumption holds, competence levels will be better than random. However the converse does not hold: if agent's levels of competence are better than random it does *not* necessarily follow that agents are more likely to vote for the correct alternative than any other alternative and it is *not* necessarily the case that the extended CJT competence assumption holds. Consider an agenda with five alternatives, where alternative k_1 is the correct alternative. The agent's probabilities for voting for each of the alternatives are as follows:

$$\Pr(vote = k_1) = 0.3$$

$$\Pr(vote = k_2) = 0.4$$

$$\Pr(vote = k_3) = 0.1$$

$$\Pr(vote = k_4) = 0.1$$

$$\Pr(vote = k_5) = 0.1$$

Since there are five alternatives on the agenda, random competence would be $\frac{1}{k} = \frac{1}{5} = 0.2$. Since the probability of voting for the correct alternative k_1 is $p = 0.3 > 0.2$, agents have better-than-random levels of competence. However since the probability of agents voting for the correct alternative is *not* greater than the probability of voting for the correct alternative (agents are more likely to vote for the incorrect alternative k_2) the competence assumption of the extended CJT does not hold.

So, better than random levels of competence are necessary but not sufficient for the extended CJT competence assumption to hold. Establishing that competence levels are better than random does not establish that the competence assumption holds[48]. To interpret the Disjunction Problem charitably, we could say that the extended competence assumption requires at least better-than-random levels of competence. If we cannot establish that agents have at least this level of competence then we cannot establish that agents are more likely to vote for the correct alternative than to vote for any other alternative.

The second point of exposition required before the Disjunction Problem can be presented is an argument for $p > 1/k$, an argument for why we can assume that agents are more likely than random to vote for the correct alternative. The argument is essentially that agents would have a random level of competence if they just guess what correct alternative is, for example if they allocate their vote by tossing a k-sided dice. If agents have the smallest amount of truth-conducive information, if they have even the slighted idea what the correct alternative is, then $p > 1/k$. It seems trivial to assume that a moderately capable agent would

[48] Except in the special case where error is distributed evenly across the incorrect alternatives. An agenda of size $k = 2$ is such a special case, since there is only one incorrect alternative and so all the error is distributed evenly on this one incorrect alternative.

be *more likely* than random to vote for the correct alternative. Let's term this argument the 'random competence' argument.

Once we have established that the extended competence assumption requires at least better-than-random levels of competence and once we have an argument for competence levels being better than random, the Disjunction Problem has its target. The crux of the Disjunction Problem is that there is no principled way to determine the number of alternatives that should be on an agenda for a social choice and therefore that there is no reason to assert that the competence assumption of the CJT holds.

The Disjunction Problem can be presented in two ways. Firstly, suppose that initially $k = 10$ i.e. our agenda is comprised of alternatives $k_1, k_2, \ldots, k_{10}$. Better-than-random competence levels would require competence $p > 1/10$. But suppose at a later stage 9 of the original 10 alternatives are joined in a disjunction so that $k' = 2$ i.e. our agenda is now k_1, k_2' where $k_2' = k_2 \vee k_3 \vee \ldots k_{10}$. With a revised agenda better-than-random levels of competence now require competence $p' > 1/2$. The two agendas (of 10 or 2 alternatives) are logically equivalent. Merely as a result of revising the way in which we describe the agenda, our assumed level of competence has increased from $p > 1/10$ to $p' > 1/2$. It may seem unremarkable that an agent would have a better than 1/10 chance of voting for the correct alternative on an agenda but merely as a result of reframing the description of the choice an agent is assumed to have a better-than 50% chance of voting for the correct alternative – quite high, given that there could be more than 10 alternatives on the agenda.

Equivalently we could start with an agenda of $k = 2$. Better-than-random levels of competence require competence $p > 1/2$. But we could represent one of these two alternatives as a disjunction of 9 alternatives, meaning $k' = 10$ i.e. our agenda is now $k_1, k'_2, \dots, k'_{10}$. Better-than-random competence now requires $p' > 1/10$. Merely as a result of revising the way in which the agenda is presented, the assumed level of competence has decreased from $p > 1/2$ to $p' > 1/10$.

The concern behind the Disjunction Problem is not trivial. Estlund cites an example of blind men and an elephant. A group of blind men are allowed to touch an animal and are then asked whether it is an elephant or not. In such a binary choice competence should be $p > 1/2$. But not being an elephant is equivalent to being a hippopotamus, or being a rhinoceros, or being a mule, or being a horse and so on. While it may initially seem obvious that an agent will be better than random at determining the correct alternative from an agenda of elephant/ not elephant, it seems implausible that they will have a better than 50% chance of correctly identifying that the animal is an elephant, given all the other possible animals it could be.

<u>The Disjunction Problem is misguided</u>

The random competence argument in favour of the CJT competence assumption holding is very weak. Firstly, establishing that agent's levels of competence are better than random does not establish that the extended CJT competence assumption holds (except in the special case where error is distributed evenly across incorrect alternatives). It is therefore difficult to see why anyone would advance the random competence argument in support of the CJT competence assumption holding.

Secondly, knowing the size of an agenda does into license anyone to infer anything about the competence level of an agent. It is too crude to place all social choice problems with the same agenda size into the same category and assume that agents will have a level of competence better-than-random. There will be some social choice problems with k alternatives on the agenda where agents have no relevant information whatsoever, where competence is $p = 1/k$ and the extended CJT competence assumption does not hold. There will also be some social choice problems with k alternatives where agents receive misleading information such that competence $p < 1/k$ and the extended CJT competence assumption does not hold. Finally for agenda size k there will be some social choice problems where agents have truth-conducive information such that competence is $p > 1/k$, where agents are more likely to vote for the correct alternative than any incorrect alternative and therefore where the extended CJT competence assumption does hold. The random competence argument for the CJT competence assumption holding is absurd since defining a reference class according to agenda size k is far too crude.

Thirdly, we can present a reductio argument *against* the random competence argument. Assume that the random competence argument justifies the CJT competence assumption holding. According to the random competence argument the competence assumption holds in all cases. We know as an empirical fact that in some social choice problems the competence assumption does not hold (because of biases in agents or misleading information received by agents). Therefore the random competence argument cannot be correct.

Given these three criticisms of the random competence argument, the emphasis that Estlund places on the random competence argument as the defence of the CJT competence assumption seems like a straw-man argument.

Estlund states "…without that assumption [the random competence argument], or some substantive support for the competence assumption, the jury theorem gets us nothing." (p.230). If the random competence argument was successful then it would have provided a sufficient (but not necessary) justification for the CJT competence assumption holding. It does not follow that if the random competence argument is defeated the competence assumption does not hold. At worst, the failure of the random competence argument simply means the competence assumption is currently without support. As Estlund acknowledges, there may be other arguments for the competence assumption holding and I will present one such argument later in this chapter.

<u>Agenda setting: the concentration of error and dispersal of 'competence'</u>

The presentation of the Disjunction Problem in terms of random competence is something of a red-herring. As shown above, it is true that if the CJT competence assumption holds then agent's level of competence is better than random. However what the Disjunction Problem *actually* highlights is how the framing of a social choice problem by a social planner can affect whether the

CJT competence assumption holds or does not hold. The CJT competence assumption may fail to hold if error is concentrated on one alternative, or if competence is dispersed across several alternatives.

The concentration of error

Suppose that as a matter of fact the animal in the next room is an elephant and the social planner fixes the description of the one correct alternative on the agenda as ‘elephant’. Whether the CJT competence assumption holds or not depends on whether agents are as likely or more likely to vote for another alternative that is incorrect. And whether agents are as likely or more likely to vote for another alternative that is incorrect in turn may depend on how many incorrect alternatives there are on the agenda.

Suppose $k = 7$. For example the agents may face the following agenda:

The animal in the next room is:

k_1 an elephant; or

k_2 a hippopotamus; or

k_3 a rhinoceros; or

k_4 a mule; or

k_5 a horse; or

k_6 a dog; or

k_7 none of the above.

Suppose that the probabilities for voting for each of the alternatives are as follows:

$$\Pr(vote = elephant) = 0.3$$

$$\Pr(vote = hippopotamus) = 0.1$$

$$\Pr(vote = rhinoceros) = 0.1$$

$$\Pr(vote = mule) = 0.1$$

$$\Pr(vote = horse) = 0.1$$

$$\Pr(vote = dog) = 0.1$$

$$\Pr(vote = none\ of\ the\ above) = 0.2$$

Here the extended CJT competence assumption holds – agents are more likely to vote for the correct alternative 'elephant' than they are to vote for any of the incorrect alternatives.

Now suppose that instead the agenda is comprised of two alternatives as follows:

The animal in the next room is:

k_1 an elephant; or

k_2 none of the above.

The competence of an agent on this revised agenda should remain at $p = 0.3$. The revised agenda is logically equivalent to the original agenda, and the revised agenda does not give the agent any more information than the original agenda, so the probabilities of an agent voting for the correct alternative on each agenda should be the same. If agent competence is $p = 0.3$ then all of the agent error $(1 - 0.3 = 0.7)$ is concentrated on one incorrect alternative as follows:

$$\Pr(vote = elehant) = 0.3$$

$$\Pr(vote = \neg elephant) = 0.7$$

Therefore it is not the case that agents are more likely to vote for the correct alternative than any incorrect alternative and so the CJT competence assumption does not hold.

The social planner, in setting an agenda, controls how the error of an agent will be distributed across the incorrect alternatives. If she restricts the number of incorrect alternatives on the agenda then she may concentrate the error of agents to such an extent that the agent is more likely to vote for the incorrect than correct alternative. There is no principled way for a social planner to set an agenda, in such a way as to avoid concentrating error on a specific incorrect alternative, to such an extent that the extended CJT competence assumption does not hold.

The dispersal of 'competence'

Suppose the social planner adjusts the description of the one correct alternative on the agenda. Whether the CJT competence assumption holds or not depends on whether agents are more likely to vote for the correct alternative than they are to vote for an incorrect alternative, and this in turn depends on how refined the description of the correct alternative is.

Suppose $k = 2$. For example, agents might be facing the following agenda:

The animal in the next room is:

k_1 an elephant; or

k_2 none of the above.

Suppose the competence of agents in this case is $p = 0.6$. Given that there are only two alternatives on the agenda, it follows that the error is concentrated on one incorrect alternative and the probability of voting for the incorrect alternative is $(1 - 0.6 = 0.4)$. Here agents are more likely to vote for the correct alternative than any other alternative and so the CJT competence assumption holds.

Suppose that the description of the correct alternative is further refined. The alternative of 'elephant' is equivalent to the alternative of 'African bush elephant or African forest elephant or Asian elephant'. Now agents face the following agenda:

The animal in the next room is:

k_1 an African bush elephant; or

k_2 an African forest elephant; or

k_3 an Asian elephant; or

k_4 none of the above.

It is an open question how the 'competence'[49] is dispersed when the correct alternative is split into more refined alternatives. Agents have a 0.6 probability of correctly identifying the alternative as an elephant. It may be the case that agents are just as able to correctly identify the animal as an African bush elephant as they are to correctly identify the animal as an elephant. In such a case, the CJT competence assumption holds. However it is more plausible that agents are *less* able to correctly identify the animal as an African bush elephant than they are to correctly identify the animal as an elephant. As such the 0.6 probability of voting for the correct alternative of 'elephant' must be dispersed across the more refined alternatives of 'African bush elephant', 'African forest elephant' and 'Asian elephant'. Where there are four alternatives on the agenda it is possible the 'competence' is distributed as follows:

$$\Pr(vote = African\ bush\ elephant) = 0.3$$

$$\Pr(vote = African\ forest\ elephant) = 0.15$$

$$\Pr(vote = Asian\ elephant) = 0.15$$

[49] The term 'competence' is presented in scare quotes since competence is defined as the probability of voting for the correct alternative. If this probability is distributed across some alternatives that are incorrect then it is somewhat misleading to describe it as competence.

Agents still have a 0.4 probability of voting for the incorrect alternative of 'none of the above' hence it is not the case that agents are more likely to vote for the correct alternative than any other alternative and the extended CJT competence assumption does not hold. There is no principled way for a social planner to set an agenda in such a way as to avoid dispersing the 'competence' of agents across alternatives to such an extent that the extended CJT competence assumption does not hold.

An open-ended agenda

The way in which a social planner sets an agenda can determine whether or not the CJT competence assumption holds. The framing of the agenda may result in agent error being concentrated on a particular alternative to such an extent that agents are more likely to vote for an incorrect than the correct alternative. Similarly, the framing of an agenda may result in a description of the correct alternative that is so refined that agents have low probability of voting for it and agents are instead more likely to vote for an incorrect alternative.

If setting an agenda in advance can mean the competence assumption of the extended CJT does not hold the social planner might choose to present agents with an open-ended agenda. However, an open-ended agenda brings with it additional problems.

Suppose agents are told that the social choice problem is to identify the animal in the next room, but they are not told what the animal might be. One by one the blind men are allowed to enter the room and touch the animal. When they leave the room the agents cast their votes for what they judge the animal to be. Suppose the five agents cast their votes as follows:

Agent i: Rhinoceros

Agent j: Rhinoceros

Agent k: African bush elephant

Agent l: African forest elephant

Agent m: Asian elephant

What is the plurality winner in this example? Or more to the point, what is the animal in the next room likely to be given the judgements expressed by agents? This seems like an open question. On one interpretation of the votes the alternative ‘rhinoceros’ is the plurality winner and so the animal in the next room is probably a rhinoceros. On a different interpretation the alternative ‘elephant’ is the plurality winner and so the animal is probably an elephant. The problem with an open-ended agenda is that we cannot objectively interpret or make use of the information agents have secured via their searches. In the absence of an agenda shared by all the agents we do not know whether agents were voting on a tacit two placed agenda of:

The animal in the next room is:

k_1 an elephant; or

k_2 not an elephant.

in which the plurality winner is ‘elephant’; or if agents were voting on a tacit two-placed agenda of:

The animal in the next room is:

k_1'a rhinoceros; or

k_2' not an rhinoceros.

in which case the plurality winner is ‘not rhinoceros’; or if agents were voting on a tacit four-placed agenda of:

The animal in the next room is:

k_1''a rhinoceros; or

k_2'' an African elephant; or

k_3'' an African bush elephant; or

k_4'' an Asian elephant.

in which case the plurality winner is ‘rhinoceros’.

Different agents may also have been voting on different tacit agendas. In the absence of a set agenda there is no objective way of counting votes for a particular alternative and no objective way of counting votes against a particular alternative.

As an aside, it is tempting to see Condorcetian mechanisms at work in ordinary choice problems. But the lack of a single, shared agenda for agents means such applications are tenuous. For example, when visiting foreign cities it is nice to sample the best of the local cuisine. A useful heuristic for finding the best local cuisine is to go to the restaurant where most of the locals seem to go. Setting aside the problem of information cascades[50] and group think[51] this heuristic seems like an instance of the CJT in action. Suppose that there are a number of restaurants in the city centre, that the local people are able to make good but imperfect judgements regarding the quality of a restaurant and that the locals make independent decisions where to dine. The extended CJT implies that as the size of the dining population increases the probability that the best restaurant has more diners than any other restaurant increases. However this reasoning requires us to interpret an agent's presence in a restaurant as their judgement that this restaurant provides the best local cuisine. Some locals may choose a restaurant on that basis, but other locals may choose a restaurant because it is the cheapest or because it has the best wine list or because it has the best view. Again, without a single shared agenda the CJT framework is simply inapplicable.

The Disjunction Problem is misguided. However the analysis of the Disjunction Problem does highlight the problem of agenda setting and we are immediately placed on the horns of a dilemma. A social planner can choose to either set an agenda in advance or not set an agenda in advance. If a social

[50] An information cascade might occur as follows. Suppose the first agent makes an independent judgement of a restaurant's quality. A second agent chooses the restaurant because they can see the first agent dining there. The third agent chooses the restaurant because they can see two other diners there, and so on. Although it may seem that a number of different diners have made independent judgments of the restaurant's quality in fact all but one of the judgements depend on one agent.

[50] Group think might occur if a number of diners actually hate the restaurant, but they stay in the restaurant because they don't want to be seen as an outsider.

planner were to set an agenda in advance there is no way in principle to avoid concentrating the error of agents on a particular incorrect alternative, or to avoid dispersing ‘competence’ across several alternatives, such that the CJT competence assumption does not hold. If the social planner does not set an agenda in advance then there is no way to objectively count the number of votes for particular alternatives and therefore there is no objective plurality winner. Neither of the options is attractive.

An argument for the CJT competence assumption holding

No one should argue that the CJT holds unconditionally. The CJT only asserts that *if* the independence and competence assumptions hold *then* the probability of a correct social choice is monotonically increasing in group size and in the limit tends towards certainty. The CJT only gives conditional support to the truth-tracking ability of democratic decision making. The CJT could only support an epistemic defence of democratic decision making if the antecedent of the conditional holds i.e. only if the competence and independence assumptions hold. Estlund is quite right to state that "...the assumption that voters are better than random is not freely available, but would need some argument" (p.231). He is right to demand a separate argument for the competence assumption holding if we want to use the CJT as a justification for an epistemic defence of democratic decision making.

What should we expect from a defence of the CJT competence assumption? We should not expect the CJT competence assumption to hold in all social choice problems. Agents may have systematic biases in particular types of social choice problem. For example, agents may have a bias in favour of the status quo; as such they may choose to re-elect a President even when the opposition candidate is superior. Even when agents are facing a type of social choice problem where they do not have systematic biases there remains a possibility that they will receive misleading information. For example, jury members may be carefully selected to avoid agents with biases. However the jury may be presented with inaccurate witness statements that lead them to wrongly convict an innocent defendant. Where agent competence is worse than random (because agents have systematic biases or because they have received misleading information) the probability of a correct plurality winner is *decreasing* in group size and in the limit tends towards *zero*[52]. If we want to employ the CJT for an epistemic defence of democracy we need to show that the competence assumption holds most of the time, in the relevant types of cases.

The original formulation of the Disjunction Problem in terms of random competence made the mistake of looking for an in-principle justification of the competence assumption holding. Similarly, the reformulation of the Disjunction Problem showed that we cannot guarantee in-principle that the social planner has framed the agenda in such a way that the competence assumption holds. But this is only a concern if we want an in-principle justification for the competence assumption holding. Given the variety of different circumstances involved in different social choice problems an in-

[52] See the presentation of the classic CJT in Grofman, B., Owen, G. and Feld, S.L. (1983) Theorem I.

principle justification of the competence assumption which applies to all social choice cases is not credible.

Instead, the best hope for justifying the claim that the CJT competence assumption holds is to identify as closely as possible the type of social choice problem where the competence assumption does hold. For example, it would be wrong to argue that the competence assumption holds in all jury trials since there is empirical evidence that juries sometimes make mistakes. Furthermore it would be wrong to argue that the CJT competence assumption holds in most jury trials since this is too coarse a reference class. Rather, it could be argued that the competence assumption holds in most jury trials where there is careful selection of jurors to avoid biases, where there are proper rules of evidence and where the police have collected sufficient evidence. We know that the CJT competence assumption tends to hold in a reference class of social choice problems such as this since very few of the verdicts are overturned on appeal. If we can show that the circumstances of a particular social choice problem are like those of the reference class of problems where the CJT competence assumption tends to hold, then we have a justification for the competence assumption holding in the particular case.

A consequence of this defence of the CJT is that the asymptotic limit to the probability of a correct social choice is not certainty but the probability that the social choice problem is non-misleading (a value greater than 0.5, but less than certainty). This issue is considered again in chapter 7.

Summary

Multiple alternatives posed a particular problem for the classic CJT and majority rule. We can vary the number of elections and size of agendas to ensure the sheer number of possible alternatives does not degrade the competence of agents too much. As such, discussions in the thesis that presume a dichotomous choice are just a rhetorical convention: it is more straightforward to discuss dichotomous choice situations and it should be taken as given that the discussions apply to cases where there are multiple alternatives.

Chapter 4: The generation of the inputs to aggregation procedures.

This chapter provides an analysis of how the judgement-generating factors of competence, independence and transparency should be interpreted and an analysis of how these form. In the process we will consider a taxonomy of the causal factors of an agent's judgement, including truth-conducive evidential and background information, that generate the variables of competence and independence and that are in turn pooled by the aggregation procedures. The taxonomy takes the distinctions from Dietrich (2008) as its starting point. However, the taxonomy in this chapter differs in two respects. Firstly, there is a matter of emphasis. It is not just the truth-conducive evidential information that is crucial in forming an agent's competence. The truth-conducive background facts, such as a good education, are just as (if not more) important in forming the competence of an agent. Secondly, the taxonomy in this chapter shows how the competence and independence relations of agents develop over time. While we are most interested in the competence levels and independence relations of agents at the time at which they cast their votes, we should also be concerned with how these variables develop over time. This is particularly important for expanding the account of how agents track the truth into a two stage process of the search for information followed by the pooling of information via aggregation procedures.

The problem

Aggregation procedures generally, and majority voting in particular, play a prominent role in existing accounts of epistemic democracy, in defences of democratic decision making on the basis of its capacity to track the truth.

The inputs to an aggregation procedure are the judgements of individual agents, recorded as votes for or against an alternative. These judgements are determined by factors such as the competence of agents, the independence of agents and transparency of agent's competencies. The three different aggregation procedures of expert dictatorship, negative reliability unanimity rule and majority rule can track the truth given certain types of judgement-generating factors.

For example the Condorcet Jury Theorem (CJT) tells us that as group size increases, majority voting will be more likely than random, more likely than an individual and likely simpliciter to track the truth if the average level of competence is greater than 1/2, the distribution of competencies in the group is symmetric about the mean and the votes of agents are independent. It does not matter whether the competence of agents is transparent, though if competence were transparent we could apply weights to the votes of agents in proportion to their competencies and increase the probability of a correct majority verdict.

The existing accounts of epistemic democracy that focus on aggregation procedures only give conditional support to the truth-tracking ability of groups. They show how groups can track the truth given certain types of judgement-generating factors. They are silent on how these judgement-generating factors are themselves generated or whether they are plausible. More particularly, the existing aggregative accounts of group truth-tracking begin at the point at which agents already have a set level of competence, in a particular distribution, with certain independence relations holding and the transparency or otherwise of competence pre-determined. But agents do not have high competence a priori; it cannot be taken as given that agents will have information regarding the correct alternative on an agenda. Nor is it the case that the required independence relations will hold a priori or that the transparency of competence is established. We need an account for how the features of a group of agents, including competence levels, transparency of competence and independence relations, develop. The truth-tracking institutional features of some aggregation procedures can provide a conditional epistemic justification for group decision making; an account of the formation of the judgement-generating factors will provide the antecedent to this conditional justification.

Providing an account of how the judgement-generating factors for aggregation procedures form will also improve our analysis of the epistemic power of the aggregation procedures. The truth-tracking ability of an aggregation procedure is due to a combination of the institutional features of the aggregation procedure and the inputs to the aggregation procedure. We can see the impact that the institutional features of the aggregation procedure have on the ability of a group to track the truth by noting that different aggregation procedures (such as expert dictatorship, negative reliability unanimity rule and majority rule) will have different probabilities of generating the correct social choice, given the same inputs (given the same group of agents with set levels of competence, independence relations and transparency or otherwise of competence). We can see the impact that the inputs to the aggregation procedure have on the ability of a group to track the truth by noting that the probability of a correct social choice will vary if we keep the aggregation procedure (like majority rule) fixed, but we change the judgement-generating factors (for example we increase the competence of agents, or we change the independence relations). In sum, the truth-tracking ability of aggregation procedures is due to both the way in which the institutional features of aggregation procedures pool the information dispersed in the judgements of individual agents (i.e. the social epistemic mechanisms) and to the amount of information contained in the judgements of the individual agents themselves.

In the following sections, I set out precisely how the judgement-generating factors of competence, independence and transparency of competence are generated. In doing so it should become clearer how the variables should be interpreted and what features of real-world social choice problems they capture. This explanation also requires that we give consideration to the ‘informational environment’ the agents face: the set of possible causes of agents’ votes, including the truth-conducive evidential and background information available to them.

Competence

The competence of an agent is defined as the probability that this individual agent votes for an alternative, given that it is correct. Formally competence conditional on the state of the world is defined as $p_i = \Pr(v_i = x|x)\,, \forall_{x\epsilon}\{1,0\}$, where p_i is the probability that agent i will vote for the correct alternative, v_i is the vote or judgement of agent i and x is the state of the world (which can take values 1 or 0). Because competence is a probability, it belongs to the interval [0,1]. All truth-tracking aggregation procedures recommend that groups contain at least some agents with high levels of competence.

Interpreting competence

The competence of an agent represents the probability of an event occurring, namely the probability that a particular agent will vote for the correct alternative. The probability captures an epistemic uncertainty an observer or modeller or *social planner* has over that event occurring. In any real social choice problem with a correct alternative (such as an election or a jury trial) each agent (each voter or juror) will either cast a vote for the correct alternative or they will cast a vote for the incorrect alternative[53]. The agent's vote for a particular alternative is determined by the combination of their causal influences. If the observer were aware of all the causal influences of an agent, all randomness in the agent's vote would disappear and the agent would vote for the correct alternative with a probability of either 0 or 1. But the observer is not aware of all the causal influences on an agent's vote, and which of the two events will actually occur (whether the agent will vote for the correct alternative or the agent will vote for the incorrect alternative) is not known in advance to the observer.

[53] Abstentions are ruled out.

Precisely how agent competence is interpreted does not matter for the formal results, but it is of philosophical interest. If we address the question of how to interpret agent competence we must also address the question of how the level of agent competence is assessed. Edelman (2002) sets our three different interpretations of how randomness enters a CJT model: the random model, pooling model and aggregation model. Each of these interpretations takes the perspective of a social planner or observer, that is someone who can 'observe' the voting behaviour of agents, and who may be in a position to make institutional decisions over which judgement aggregation procedure to employ. Under the random model, the votes or judgements really are like the tosses of a coin: to say that an agent has a competence level of $p_1 = 0.6$ is to say that there is some objective randomness in the agent's vote. As Edelman notes, if we accept the random model, then if the election was repeated and an agent were to cast their vote again on the same agenda the agent may well vote differently this time. The interpretation of agent competence as an objectively random process seems inapplicable to the cases addressed in this thesis. If agents are faced with an agenda with one correct alternative, if agents have gathered evidence to inform their judgement and then cast their vote in line with what they honestly believe to be the true state of the world, then we would expect agents to vote in exactly the same way every time the elections is re-run. As the random model of competence is irrelevant for our purposes, we can ignore the question of how agent competence levels are determined under this model[54].

[54] Estlund (2008) and List and Goodin (2001) both explain the CJT with reference coin tosses, but I take it that these explanations are intended only as analogies. The coin toss examples are ideal for explaining the law of large numbers, which underlies the CJT, even if it is implausible to think of human agents as objectively random devices.

According to the polling model, the particular problem a group faces is held fixed but there is uncertainty over which individuals will comprise the group. To say that agent competence is $p = 0.6$ is to say that 60% of the wider population will cast their votes for the correct alternative (and 40% of the population will vote for the incorrect alternative). If we take random samples from the wider population to form the voting group then there is a 60% chance that a given agent will vote correctly. As Edelman notes, this model makes no assumption that the voters have any information whatsoever about the true state of the world which generates their judgements. The polling model might be appropriate for non-epistemic social choices, for example where a group needs a collective decision over whether to prioritise education or health spending and agents merely express their preferences. However, in the epistemic social choice problems considered in this thesis the polling model is inapplicable and we can again put to one side the question of how agent competence levels are determined.

According to the aggregation model, the composition of the group is held fixed, but there is some uncertainty over the particular problem the group will face. To say an agent has a competence level of $p_i = 0.6$ is to say that, of all the problems in a suitable reference class that agent faces, the agent gets 60% of them right. If it is possible to tell what the true state of the world is (after agents have cast their votes and independently of the social choice outcome) then it is possible to determine the competence level of agents. Here a social planner need only determine the long range success of a particular agent at casting correct votes to determine that agent's competence level. For a judge on a panel, for example, we might determine their competence level by determine the proportion of their judgements that were overturned on appeal.

The interpretation of agent competence that I advance, that of the subjective assessment of a social planner, is consistent with Edelman's aggregation model. Furthermore the aggregation model of agent competence seems most relevant to the epistemic setting in which the CJT is applied. However, the aggregation model requires some expansion. The aggregation model as currently presented does not, as Edelman claims, provide an account of how the information dispersed in the judgements of individual agents is pooled or aggregated into the social choice. Just as in the pooling model, it is conceivable with Edelman's interpretation of the aggregation model that the votes of agents are determined by entirely non-informational causes. The fact that an agent votes correctly in 60% of cases does not mean that the agent has received truth-conducive information that makes them vote for the correct alternative in 60% of cases, and misleading information that means the agent votes incorrectly in the remaining 40% of cases. It could be that the agent casts their vote exclusively in line with their preferences (their preferred alternative happens to coincide with the correct alternative 60% of the time).

By considering how the social planner might determine the level of agent competence, we can see how judgement aggregation procedures such as majority rule can actually operate as information pooling mechanisms. A social planner might determine the competence level of agents by assessing the agent's long range frequency of voting for the correct alternative in a suitable reference class of problems. However with this approach we face the problem of determining what a suitable reference class is. For example, how do we assess the competence level of a particular judge on a panel that is about to consider an important murder case? Do we look at the proportion of their career decisions upheld? Or do we look at the proportion of their *recent* decisions upheld? Or do we consider the proportion of their decisions on murder trials that were upheld? Whichever approach we take, there remains the risk that the next murder case the judge faces is nothing like the previous cases they have faced and so the judge's previous performance is not a reliable indicator of future performance.

There are two further ways (other than assessing the long range performance of agents) in which a social planner might assess the competence level of agents. Ladha (1992) explicitly talks of majority rule as a mechanism to "…assimilate decentralised information about the alternatives." (p.619). In Ladha's example, a group has to decide whether the bias on a coin is such that the probability of heads is 0.6 or 0.3 (one of these is the true state of the world). Each agent privately observes a certain number of tosses of the coin before casting their judgement as to the bias of the coin. Edelman (interpreting Ladha) states that "…his [Ladha's] description is essentially that of my aggregation model in which the issues correspond to the private information gotten by each voter from the flipping of the coins." (p.335). However, Edelman's interpretation is misleading. The private set of coin tosses observed by each agent means that each agent receives *a different set of evidence*. Each agent faces the *same issue*, namely determining the bias on the coin. A more accurate interpretation of

Ladha's example is that the social planner will see the set of coin tosses that each agent observes. But neither the agent nor the social planner can predict whether the set of coin tosses witnessed by the agent will be representative of the coin bias or not[55]. The randomness enters the model, not because the make-up of the group is uncertain or the issues the group faces is uncertain, rather it is uncertain whether *the evidence itself* is misleading or truth-conducive. Nevertheless, this interpretation of competence does account for why majority rule (among other judgement aggregation procedures) act as information pooling mechanism. Both the agent and the social planner are aware of all the information that generates the judgements of the agents (the private sets of coin tosses), and which is then pooled in the social choice.

A second way in which a social planner might assess the competence level of agents is if they observe the evidence that an agent receives, but are uncertain how this will influence the way in which the agent votes. This uncertainty might occur because the social planner is unaware of the other (non-evidential) factors influencing an agent's vote, or because the social planner is aware of all the factors influencing the agent's vote but is unaware of how these factors interact in an agent's internal deliberation. This interpretation of the randomness in an agent's vote seems most natural to apply to the social choice problems that are the concern of this thesis. For example, a detective might know that a defendant is guilty of murder. The detective knows the evidence that was put to the jury. However the detective does not know how the jury will

[55] Again, by the law of large numbers, if the set of coin tosses observed by the agent is large then the average number of heads will tend to be very close to the actual bias of the coin. But if the set of coin tosses observed by the agent is small, it is less likely to be close to the actual bias of the coin.

interpret the evidence and how compelling the evidence will be in determining the jurors' judgements. Similarly, a political pundit may know which presidential candidate will create the most jobs[56]. The pundit may also know (or at least have some idea of) the information voters have received about the candidates. However, the pundit does not know how the voters will make use of this evidence when deciding how to vote. Again, because the social planner in these examples can see some of the information dispersed among the different agents in the group it makes sense to talk of judgement aggregation procedures pooling this information into the social choice.

If the aggregative model of how to interpret agent competence is expanded to allow for the social planner to be aware of some but not all of the information influencing an agent's vote, then the secondary question of how the social planner determines the value of an agent's competence becomes more important. On the one hand it seems entirely plausible to suggest that a social planner could assess the competence of an agent as $p_1 = 0.6$ if the social planner can see the evidence the agent has received. There is subjective randomness in the votes of agents (competence is not 1 or 0) precisely because the social planner is not certain how the agent will vote. On the other hand, different social planners witnessing the same agent receiving the same piece of evidence may come to different conclusions as to what the competence of the agent is. This issue is significant and too broad to address here. However, I will gesture in the direction of a solution. As models of consensus formation such as the Lehrer - Wagner model imply, provided that the social planners can

[56] Let's suppose, for the sake of the argument, that the objectively correct alternative in a presidential election is the candidate who will create the most jobs.

share their judgements as to an agent's level of competence, and provided that each social planner respects the view of every social planner then eventually all social planners will agree on what the level of an agent's competence is given the information the agent has received.

If an agent is to vote for the *correct* alternative the agent must have some causal factor which makes them vote for the correct alternative. For the observer to be justified in the assumption that the competence of an agent is $0.5 < p_i < 1$, the social planner must be aware of some of the causal factors that influence an agent towards voting for the correct alternative, but not be aware of all the causal factors influencing an agent's vote.

<u>A taxonomy of causal factors</u>

The causal factors influencing an agent's vote will be many and varied. An agent may vote for a particular alternative for non-cognitive reasons, for example they may just have a gut instinct that a particular alternative is correct. Environmental factors may also have a causal influence on votes. For example, poor lighting may make it difficult for a voter to read their ballot paper. Losing the World Cup may make the voter more pessimistic when they cast their vote. The background of an agent may have an influence on an agent's judgement. For example, an agent with degrees in mathematics or science will develop skills that mean they are very likely to vote for the correct alternative in a mathematical problem. An agent who lacks a formal education in mathematics would be less likely to vote correctly in such a problem.

Often the causal factors influencing an agent's vote will be informational in nature. Casting a vote involves expressing a judgement as to the correct alternative. This is an inherently cognitive activity which will be influenced by, inter alia, what the agent has seen and read and discussed with their friends, combined with what they have learned over the course of their life.

It will help to develop a taxonomy of the types of causal factors that influence an agent's vote. We can take as a foundation the taxonomy provided in Dietrich (2008) where distinctions are drawn between the evidential/ non-evidential, common/ private and truth-conducive/ misleading factors influencing agents' votes. In particular I want to place emphasis on the importance of non-evidential, private, truth-conducive causal factors in forming the competence of agents while still preserving some notion of independence. I also want to emphasise that while the competence of agents at the time they cast their votes is of primary concern, we may also be interested in how the competence of agents develops over time.

Evidential/ non-evidential factors

It should be relatively easy to draw a line around a given set of evidential factors. They are "…generally observable facts that support the correctness of an alternative including the specific nature of the alternative ...and several observable events" (Dietrich, 2008, p.4). The evidential factors can also be interpreted as the indirect causal relatives of the state of the world. So in a jury trial, if the true state of the world is that the defendant is guilty of murder, the evidence such as fingerprints left at the crime scene, the DNA evidence and the witness statements are all indirect causal descendants of the act of murder and they all indicate which alternative (guilt/ innocence) is correct. In addition, the receipt for the purchase of the murder weapon is also a causal relative of the state of the world, given that the purchase of the weapon was one of the causal factors leading to the act of murder[57]. Note that evidential signals are not *direct* causal descendants of the state of the world, because no agent has direct contact with the state of the world. An agent's contact with the state of the world is mediated via chains of causes. For example, suppose the defendant placed their fingers on the knife they used in the murder and dropped it as they left the scene. A police officer later found the knife and carefully placed it in a bag. It was taken to a crime lab where a technician carefully dusted the object and lifted off a complete image of a fingerprint. This evidence was then compared with a background database of fingerprints to find a match with the suspect. This information was finally placed before the jury at the trial. Although a juror has access to this piece of evidence, which was caused by the state of the world (the act of murder), the evidence does not provide a direct causal link to the state of the world as the juror only receives the evidence via a long chain of causes.

Non-evidential factors are any causal factors on an agent's vote that are not evidential. Non-evidential causal factors carry no information as to which alternative is correct and are not causal descendants of the state of the world. "One may regard non-evidential circumstances as factors that affect whether

[57] It is important to note that I class causal relatives, and not just causal descendants, of the state of the world as evidential factors. The purchase of the murder weapon was a cause of the act of murder and not a causal consequence of the act of murder. But information regarding the purchase of the murder weapon does indicate what the true state of the world (guilt or innocence) is.

voters observe evidential circumstances and how they interpret them." (Dietrich, 2008, p.4). Although non-evidential factors carry no information on which alternative is correct, although they are not generally observable facts that support the correctness of an alternative, they do affect agents in their voting behaviour. For example, the education of jurors should not be considered as part of the evidential circumstances. The fact that a juror happens to have a biochemistry degree makes it neither more nor less likely that the defendant is guilty, and the fact that a juror has a degree in biochemistry was not caused by the fact that the defendant is guilty. However, if the trial includes evidence about DNA traces left at the crime scene, then the fact that a juror happens to have a biochemistry degree means that they are more likely to understand the evidence and vote for the correct alternative. It is in this sense that the non-evidential factors may affect the way in which an agent interprets evidential factors and in doing so influence the way they vote. Terming these factors 'non-evidential' factors downplays the significant causal influence they can have. Instead I will term the factors 'background' factors to distinguish them from the evidential factors that are causal descendants of the state of the world. Background factors include the education of an agent, their life experiences that affect their decision making, and more general propositions they learn which help them interpret evidence. Background factors include any non-evidential factor that has a causal influence on an agent's vote.

The distinction between evidential and background causal factors is important for two reasons. Firstly, agents need both types of causal factors if their competence is to be sufficiently high. Secondly, there may be different amounts of evidential and background causes available.

Agents need both background information and evidential signals for their competence to be greater than 0.5 in a dichotomous choice. Evidential information includes the nature of the agenda agents are facing, for example whether the agents are asked to vote 'guilty/ not guilty' or 'elephant/ not elephant'. If agents do not have this basic piece of evidence and instead are voting on an agenda of 'x/¬x' or '1/0' then they have no indication which alternative is correct and the probability that they will vote for the correct alternative will be 0.5. If agents do have at least some evidential information, such as the nature of the agenda, they still need at least some appropriate background factors for them to make use of the evidence. For example, an agent who does not understand that 'innocent' is a synonym for 'not guilty', or an agent who does not know what an elephant is would also only have a 0.5 probability of voting for the correct alternative even if they were told the content of the agenda. In the absence of background information agents will be unable to properly interpret the evidence. In the absence of evidential signals the background information is of no use in identifying the correct state of the world.

For the group as a whole, having a large number of diverse evidential signals and having a large amount of diverse background information are both important. The larger the quantity of evidence and the more varied the evidence available to the group, the easier it will be for the group to identify the correct state of the world. The larger the quantity and more varied the background information in a group, the better the group will interpret the evidential information. Ceteris paribus it is epistemically virtuous to include as many different agents in a group as feasible; both because of the evidential information they can contribute to the group and because of the background factors they can contribute. These claims hold provided of course that this information and these factors are non-misleading.

There may be social choice problems where the amount of evidence is limited. For example, in a criminal trial there may only be a small handful of witnesses to an act of murder. In a Presidential election, the amount of information on which voters can judge candidates could be limited to manifesto documents, official biographies, and the content of speeches and debates made during the campaign. In these cases we rely on the larger pool of background factors to appropriately interpret the limited amount of evidence. Similarly there may be social choice problems where the amount of background information is limited but the amount of evidential information may be more substantial. For example, a group of climate scientists may all use similar techniques and have similar training. Including more scientists in the group is only likely to improve their understanding of climate change if these new scientists are able to make new observations, if they are able to obtain extra pieces of evidence.

We can consider the competence of agents in certain reference classes of social choice problems (including a reference class with just one specific social choice problem). The reference class of problems could be drawn narrowly or widely, and the competence of different kinds of agents will vary according to reference class. There may be some agents who have a consistently high level of competence in a narrow reference class of social choice problems. For example, a homicide detective may have an extremely high level of competence at judging whether someone is guilty of murder, and may also have a high level of competence in the slightly broader reference class of criminal cases generally. However this detective may have a much lower level of competence at judging who the best Presidential candidate is. There may also be agents who have a consistent level of competence over quite a broad range of social choice problems. For example, a five year old child may have a consistently low level of competence across a wide range of social choice problems. Similarly an individual who has graduate degrees in both physics and moral philosophy may have such a comprehensive level of education that they have a consistently high level of competence across a very broad range of social choice problems. In each of these cases the extent to which an agent will have high or low levels of competence in a narrow or broad reference class of social choice problem is largely due to the agent's background causes. Having a broad range of background information (or background factors more generally) means that an individual is able to correctly interpret evidential information in a broad reference class of social choice problems.

We can represent the causal influences on an agents' vote in diagrams[58]. Note the causes should be interpreted as instantiations of random variables. Figure

[58] Here I follow the same format for causal diagrams employed in Dietrich and Spiekermann (unpublished a, b) except that here the causes should be interpreted as instantiations of random variables.

4.1 shows all the causes of an agent $1's$ vote and so here the vote of agent 1 will be deterministic[59]: the agent will vote for the correct alternative with either probability 1 or probability 0, conditional on the causes. In these figures, x is the state of the world, c^e is an evidential cause and c^b is a background cause. The direction of cause is represented with arrows.

Figure 4.1: an example of a complete causal network.

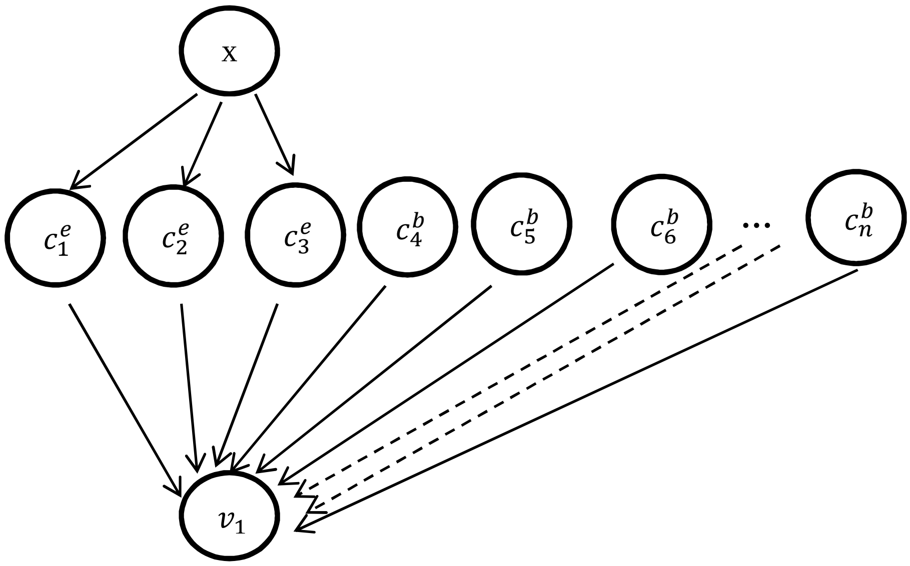

Figure 4.2 represents the causal influences on the vote of agent 1 from the perspective of an observer or social planner who is aware of (and includes) some but not all of the causes. Here there is some randomness in the vote of agent 1.

[59] Assuming there are no objectively random causes of an agent's vote.

Figure 4.2: an example of a causal network, from the perspective of an observer.

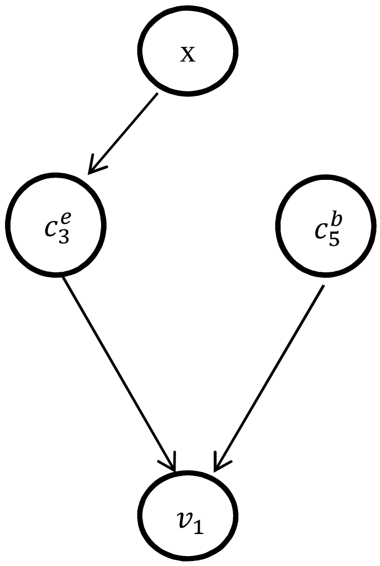

Truth-conducive and misleading causal factors

Any causal factor, be it evidential or background, can either be misleading or non-misleading. This distinction comes from Dietrich (2008). It is important to stress what is meant in this thesis by these terms. A misleading factor is one that will tend to make an agent vote incorrectly (decrease their competence). A non-misleading or truth-conducive factor is one that will tend to make an agent vote correctly (increase their competence). Misleading factors can be evidential (such as planted DNA evidence) or non-evidential (such a head cold souring a juror's mood). Non-misleading/ truth-conducive factors can also be evidential (such as actual DNA evidence left at the crime scene by the perpetrator) or non-evidential (such as a juror's degree in biochemistry). A summary of the taxonomy of causal factors influencing an agent's vote is provided in the figure below:

Figure 4.3: a summary of the taxonomy of causal factors influencing an agent's vote.

Evidential signals	Truth-conducive
	Misleading
Background factors	Truth-conducive
	Misleading

<u>Acquiring causal factors, and updating competence</u>

No evidential signal and very little in the way of non-evidential background factors will be possessed by agents a priori. Agents cannot have information about the state of the world without having contact with it. Agents may have some non-evidential information innately, but arguably most of the background information agents possess comes from them learning over time.

Over time agents obtain non-case-specific background factors, either through formal education or more generally from their experiences over their lifetime, which can influence the way they make later decisions. We can define the prior competence of agents, p_i^0, as the probability of an agent voting for the correct alternative given their current set of background factors but in the absence of any evidential factors whatsoever. This represents the competence of an agent at a time before they have considered the specific social choice problem and more particularly before they have discovered the content of the agenda. If we accept that in the absence of any evidential factors whatsoever[60] agents are

[60] Including the absence of even the details of the agenda agents are facing. Here we are applying the Principle of Insufficient Reason.

neither more nor less likely to vote for the correct alternative, then it is plausible that the prior competence of an agent in a dichotomous choice situation is $p_i^0 = 0.5$.

The competence of agents will vary as they receive new causal factors. We can term the competence of agents after they receive at least some evidential signal (such as the details of the agenda) the agent's posterior competence and we can apply a time index to the competence of agents. For example the competence of an agent i at time 1,2… t is $p_i^1, p_i^2, \dots, p_i^t$. Formally we can state the posterior competence of an agent i at time 1 as $p_i^1 = \Pr(v_i = x | x, c_1^b, c_2^e)$, where c_1^b are the prior background factors of agent i, and c_2^e is an evidential cause such as the content of the agenda[61]. Agents can receive additional information (or more generally can be influenced by additional causal factors) once they have received an initial evidential signal (such as the content of the agenda). For example, if an agent i subsequently received a further background cause c_2^b then their revised competence would be $p_i^2 = \Pr(v_i = x | x, b_i, c_1^e, c_2^b)$. If the agent then received a further evidential signal c_3^e then their revised competence would be $p_i^3 = \Pr(v_i = x | x, b_i, c_1^e, c_2^b, c_3^e)$.

The more causal factors that are conditionalised on, the less randomness there will be in the vote of the agent, conditional on these factors and the state of the world. In the limit, if all causal factors of an agent's vote are conditionalised on there will be no randomness in the vote of the agent and he or she will vote for the correct alternative with probability 1 or 0.

[61] p_1^1 is also conditional on the state of the world.

Consider two examples of how agents' competencies will vary over time. Firstly suppose the social choice involves choosing the next Presidential candidate, where the correct alternative is the candidate who will be best at managing the economy. At time t_1 an agent may be told the date of the Presidential election and be told who the Republican and Democratic candidates are. At time t_1 the little evidential information agent 1 has, combined with their prior background information, means they will have a posterior competence strictly greater[62] or less than 0.5. Subsequently (at time t_2) agent 1 watches a Presidential debate where a candidate claims that lowering taxes for the wealthy will stimulate economic growth. This is an evidential signal; a piece of information which indicates which of the candidates would be best at managing the economy. The agent who receives this signal could subsequently (at time t_3) consult economic textbooks to see whether the statement of the candidate withstands scrutiny. The economic literature assessing the impact of tax cuts on economic growth is a background factor; it does not directly imply which is the better Presidential candidate on the agenda, however an agent who gains this background information will be better able to interpret the evidential signal from the debate. If the economic literature does imply that cutting taxes increases growth then this indicates the candidate does have some economic proficiency and so is likely to be the best candidate. Alternatively, if the literature implies that tax cuts do not increase growth, then the candidate either

[62] $p_1^1 > 0.5$ if the information regarding the agenda is accurate and agent i has truth-conducive prior background information, for example, that Democratic candidates tend to manage the economy best. $p_i^1 < 0.5$ if the agenda information is accurate but the agent has misleading prior background information, for example, that Republican candidates tend to manage the economy best.

doesn't understand economics or is dishonest. Either way, the fact that the agent sought out and received extra background factors means their competence, their probability of voting for the correct alternative, has increased in this case[63].

Similarly in a jury trial, the jurors already know the defendant is charged with murder. Given their prior background factors this evidential factor means their posterior competence at time t_1 will be strictly greater or less than 0.5. For example, juror 1 may already have the truth-conducive background knowledge that most murder suspects are guilty and this factor, combined with the evidence that the agent is on the jury for a murder trial, means juror 1's initial posterior competence will be $p_1^1 > 0.5$. A different juror 2 may have misleading background experiences that lead them to distrust the police. As such their background factor, combined with the evidence that they are sitting on a jury trial, means their initial posterior competence will be $p_2^1 < 0.5$. The background experiences of juror 2 mean they are less likely to vote for the correct verdict. Later in this trial the prosecutor may introduce a new piece of evidence, such as the fact that the fingerprints of the suspect were found at the crime scene. The defence lawyer may also introduce more background information, such as the testimony of an expert witness who argues that fingerprint evidence is misleading. The competence of these agents (the jurors) will vary over time as they receive more evidential information and more background information.

[63] For example, hearing a candidate state that lowering taxes for the wealthy will stimulate economic growth may be a misleading piece of evidence i.e. it increases the likelihood an agents will vote for the incorrect candidate (the candidate who makes the statement). However, receiving the background information from the economic textbook allows an agent to see that the statement from the candidate is misleading.

As noted above, causal factors, be they evidential factors or background factors, can be misleading or non-misleading/ truth-conducive. The time-indexing of agent competence levels allows us to articulate more clearly the misleadingness or truth-conduciveness of causal factors. A causal factor c_{t+1} received at time $t+1$ is misleading if $p_i^t > p_i^{t+1}$. For example, if jurors hear an expert witness who wrongly suggests that DNA evidence is infallible, their competence may decrease. A causal factor is truth-conducive if $p_i^t < p_i^{t+1}$. For example, if jurors are presented with more true evidence that additional fingerprints were left by the defendant near the crime scene then this may increase agent competence levels.

How agents update their competence, how they translate their prior competence into posterior competence given background and evidential causal factors, is open to debate. It seems plausible that the strength of a causal factor (and whether it is truth-conducive or misleading) will depend on the combination of information an agent has, both in terms evidential signals and background information. For example consider again two agents, a lay person with a basic science education and a professor of chemistry, who receive the same evidential signal regarding climate change. If the evidential signal points to the correct state of the world, then the background information of the professor of chemistry (all that they have learnt during their education and years of research) allow them to extract significant gains in competence out of the evidential signal, meaning their posterior competence may be high. The lay person may have a much lower posterior competence given the same evidential signal because they lack the background factors required to interpret the evidential signal appropriately.

It also seems plausible that signals have decreasing marginal contributions to competence. So if agents have competence close to 0.5, any causal factor will have a significant impact on that agent's posterior competence. If competence is closer to 0 or 1 then further additional signals may have less of an impact. For example, the first witness a juror hears may convince them to vote (correctly) for guilt. The 41st witness may have less of an impact on a juror.

Although the competence of agents will vary over time, it is the competence of agents at the time when they cast their votes that matters for the aggregation procedures. However the model of the competence of agents, indexed to time, is important for two reasons. Firstly it is more realistic than the static model of competence implied by current aggregation accounts. Secondly, it will be important later in the thesis where I provide a model of how agents search for the evidential and background factors to generate their competence.

The distribution of competencies in a group depends on the competencies that different agents develop, which in turn depends on the evidential and background factors that individual agents receive. If we can account for how different agents obtain evidential and background information, we can account for the final distribution of competencies within the group.

Independence

The independence of agents, the probability that an agent will vote for the correct alternative given the vote of another agent, is important for some aggregation procedures. Informally, independence relations capture the extent to which agents will tend to vote in the same way or tend to vote differently. Formally, agent i is independent of agent j if $p_i = \Pr(v_i = x|x) = \Pr(v_i = x|x, v_j = 1)$ and $p_i = \Pr(v_i = x|x) = \Pr(v_i = x|x, v_j = 0)$. If agents are independent then the fact that one agent votes a certain way makes it neither more nor less likely that a second agent will vote for the correct alternative, given the state of the world. I have argued that the probability that an agent votes for the correct alternative is determined by the evidential signals they receive combined with the background factors they have. If two agents share at least some background or evidential factors then independence will not hold, conditional just on the state of the world. For example if the shared factors are truth-conducive then the fact that one agent votes correctly increases the probability that the second agent votes correctly i.e. $\Pr(v_i = x|x) < \Pr\left(v_i = x \middle| x, v_j = \begin{Bmatrix}1\\0\end{Bmatrix}\right)$. Examples where shared evidential or background causal factors will impact on independence relations are shown in figure 4.14 a, b and c. Common or shared factors are shaded grey[64]. In figure 4.14a agents share an evidential factor. In figure 4.14b agents share a background factor. In figure 4.14c agents share both evidential and background causal factors.

[64] This again is in line with the convention employed in Dietrich and Spiekermann (unpublished a, b). However in my diagrams the causes should be interpreted as instantiations of random variables.

Figure 4.4: examples of causal networks where agents share causal factors.

(a)

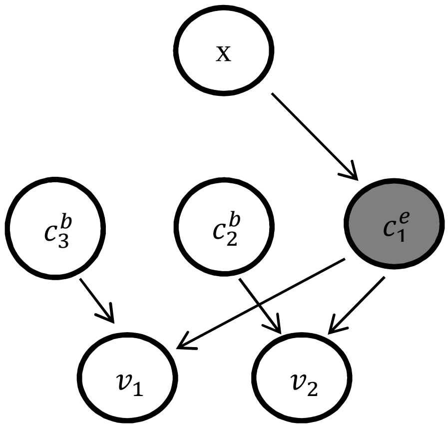

(b)

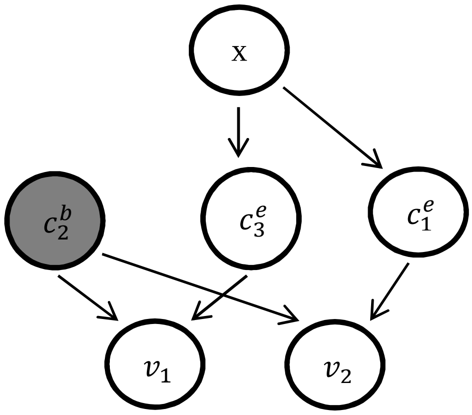

(c)

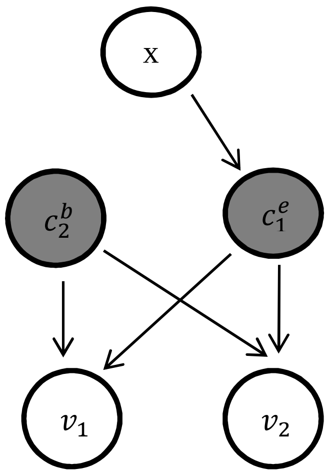

Independence can be secured, even if agents share some evidential signals or background information, if we conditionalise on all causal factors held in common. So in our three examples independence is secured as follows:

a) $$\Pr(v_1 = x|x, c_1^e) = \Pr(v_1 = x|x, c_1^e, v_2 = \begin{cases}1\\0\end{cases});$$

$$\Pr(v_2 = x|x, c_1^e) = \Pr(v_2 = x|x, c_1^e, v_1 = \begin{cases}1\\0\end{cases})$$

b) $$\Pr\big(v_1 = x\big|x, c_2^b\big) = \Pr(v_1 = x|x, c_2^b, v_2 = \begin{cases}1\\0\end{cases});$$

$$\Pr\big(v_2 = x\big|x, c_2^b\big) = \Pr(v_2 = x|x, c_2^b, v_1 = \begin{cases}1\\0\end{cases})$$

c) $$\Pr\big(v_1 = x\big|x, c_1^e, c_2^b\big) = \Pr(v_1 = x|x, c_1^e, c_2^b, v_2 = \begin{cases}1\\0\end{cases});$$

$$\Pr\big(v_2 = x\big|x, c_1^e, c_2^b\big) = \Pr(v_2 = x|x, c_1^e, c_2^b, v_1 = \begin{cases}1\\0\end{cases})$$

If independence holds after we conditionalise on common factors then there remains a concern that conditionalising removes all randomness from the subsequent votes of agents[65]. The probability of agent i voting correctly conditional just on the state of the world may be greater than 0 and less than 1 i.e. $0 < \Pr(v_i = x|x) < 1$, but the probability of agent i voting for the correct alternative given the common factors may be either 0 or 1 i.e. $\Pr\big(v_i = x\big|x, c_{i,j}^e\big) = \begin{cases}0\\1\end{cases}$. By conditionalising on factors we capture all the causal influences on an agent's vote and remove all subjective uncertainty as to how an

[65] A concern noted by Dietrich (2008).

agent will vote. This lack of randomness in the conditional competence of agents poses a problem for some of the aggregation procedures. In some aggregation procedures it is epistemically desirable for agents to vote differently. For example, in majority voting if one agent votes for the wrong alternative it is desirable that another agent votes for the correct alternative. If the conditional probability of the second agent voting correctly given the incorrect vote of the first agent is $p_j = \Pr\left(v_j = x \middle| x, v_i\right) = 0$, then the probability of a correct majority verdict will not change as group size increases. For agents to be conditionally independent, while retaining some randomness in their votes, the agents must have some evidential signals or background information held uniquely or privately[66] by them, which is not conditionalised on[67]. In the examples shown in figure 4.14, the votes of agents in (a) and (b) retain some randomness after conditionalising on common factors because we are aware that each agent has some private causal factors that are not conditionalised on. But in 4.14c, after the common factors are conditionalised on there may be no randomness in the votes of agents since the agents may not hold any causal factors privately.

We can summarise the revised taxonomy of information in the figure below:

[66] Again the distinction between public and private factors comes from Dietrich (2008).

[67] The interplay between competence and independence is addressed again in chapter 7.

Figure 4.5: a revised taxonomy of causal factors.

	Public/ common	**Private**
Evidential signals	Truth-conducive	Truth-conducive
	Misleading	Misleading
Background factors	Truth-conducive	Truth-conducive
	Misleading	Misleading

It seems implausible that in real world social choice problems every agent would have identical background factors and identical evidential signals, and so the votes of agents will always be independent conditional on the common factors[68]. Within a group of agents there may be small clusters of agents with nearly identical factors. For example, a subgroup of agents who went to the same school, studied the same subjects at university and entered the same profession will have very similar background factors and as such will interpret evidential signals in similar way. The selection of group members from the wider population needs to be careful not to capture clusters of similar agents. If, for example, a jury is comprised of individuals with the same background then they will do no better than a single juror at correctly interpreting the evidence presented to them and arriving at the correct verdict.

[68] Not even identical twins raised in the same household have identical vote-determining causal factors, since they will have at least some different experiences during their life which will impact on their votes.

Transparency

Following the definition provided in chapter 2, competence is transparent if every agent knows the competence of every other agent and they know that they know the competencies. Competence is opaque if agents do not know the competencies of other agents and they know that they do not know the competencies. We could also interpret transparency from the perspective of the social planner or observer. Competence is transparent if the social planner can see the level of competence of all the agents, and competence is opaque if the social planner cannot see the level of competence of all the agents. The transparency of competence will be a contingent matter: in some social choice problems competence will be transparent; in other social choice problems competence will be opaque. Here I present three conditions under which competence will plausibly be transparent. Firstly, there may be cases of 'zero-knowledge proofs', where agents can communicate their level of competence without sharing evidential signals. For example I can prove that I know the phone number of a friend simply by giving the friend a call. I don't have to pass on my evidence (the friend's phone number) in order to prove my competence.

Secondly, competence may be transparent because an agent is aware of the specific evidential and background factors that generate their competence and is able to show these signals to other agents. Alternatively, competence may be transparent because a social planner is aware of the specific evidential and background factors that generate agent competencies. For example, an agent may have a high competence at judging the time of their flight departure because they have a ticket with the departure time printed on it. They can show this ticket to their friends to prove they know the departure time. Note that this type of transparency can involve the sharing of the evidential signals between agents, with a corresponding impact on independence relations. This associated impact on independence may or may not have an impact on the epistemic performance of the group depending on the aggregation procedure employed.

Finally, competence may be transparent when agents are able to establish their competence via their long range success at selecting the correct alternatives in a suitable reference class of social choice problems. For example, an agent might have made the correct prediction in 9/10 of recent national elections, and as such would be expected to be highly competent at predicting the winner of the next election. The competent agent can communicate their competence without sharing their evidential signal, in that they can provide proof of their previous success without showing on what grounds they will make their next prediction. However this method of communicating competence is not completely reliable. There is a risk that future elections are not like previous elections and so the previous performance of agents might provide no justification for their future performance.

The next two chapters will consider search procedures by which groups of agents find vote-determining causal factors. We might use the notion of a search procedure here to draw a distinction between the second two notions of transparency of competence. If competence is transparent because agents can show the causal factors generating their level of competence then this is ex post (after search) transparency of competence. However if competence is transparent on the basis of an agent's long range success at selecting the correct alternatives in a suitable reference class of social choice problems, then we are employing an et ante (prior to search) notion of agent competence. Here the competence of an agent is the expected ex post competence of an agent, formally: $p_i^t = \sum_{c_t} \Pr(c_t\ )(\Pr(v_i = x|x, c_1^b, c_t\ )$, where c_t is a causal factor that may or may not be discovered by an agent.

Importantly, if competence is transparent, then agents (or a social planner) will be able to select the aggregation procedure that is optimal at tracking the truth. If competence is transparent, if the evidential signals can be shared, and if the background information of agents is roughly equivalent, then agents should all be in agreement as to the correct alternative (see Bradley 2006). If instead evidential signals cannot be shared but competence is none the less transparent then the group is still able to make institutional decisions to maximise the probability of a correct verdict. If, for example, many agents have competence greater than 1/2, weighted majority rule will be the optimal aggregation procedure. If, on the other hand, only one agent has high competence the optimal aggregation procedure is to make the high-competence agent the expert dictator[69].

Now that we have an explanation of how the causal factors received by agents generate their competence, independence relations and transparency, I will set out some examples of causal networks to show the type of judgement-

[69] Strictly speaking weighted majority rule is equivalent to expert dictatorship where only one agent has $p_i \neq 0.5$.

generating factors they produce and the type of judgement aggregation procedure that would be appropriate.

Examples of causal networks and their impact on competence and independence

Note that in the following examples we will assume that all causes (evidential and background) are truth-conducive/ non-misleading. Unless stated otherwise the diagrams do not represent all the causes, only those causes the observer or social planner is aware of. The causes should be interpreted as instantiations of a generating random variable.

In figure 4.6 below we have the perfect setting for majority rule. Each agent (represented by the votes v_1, v_2, v_3) has an evidential cause and a background cause. For example, voter 1 receives evidential factor C_1^e and background factor C_5^b. Each cause is private; no cause is shared by agents.

Figure 4.6: a causal network where each agent has private background and evidential factors.

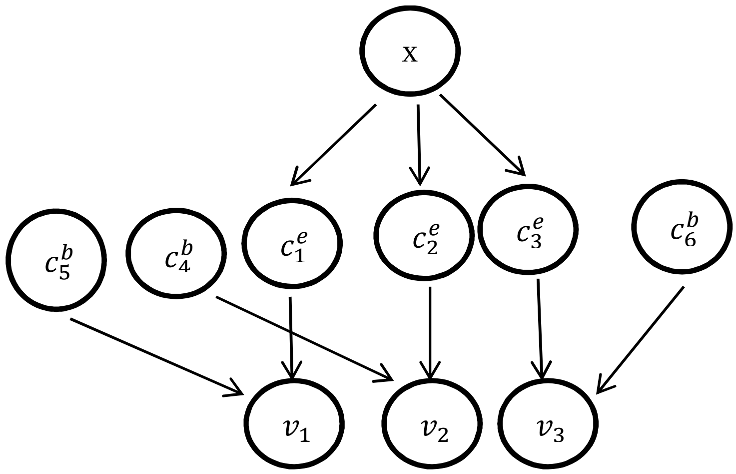

In figure 4.7 there is only one evidential cause which is shared by all agents. However, each agent has a private background cause with which to interpret the evidence, so the majority rule would still be an appropriate aggregation procedure. This example models a jury trial.

Figure 4.7: a causal network with common evidence, but private background factors.

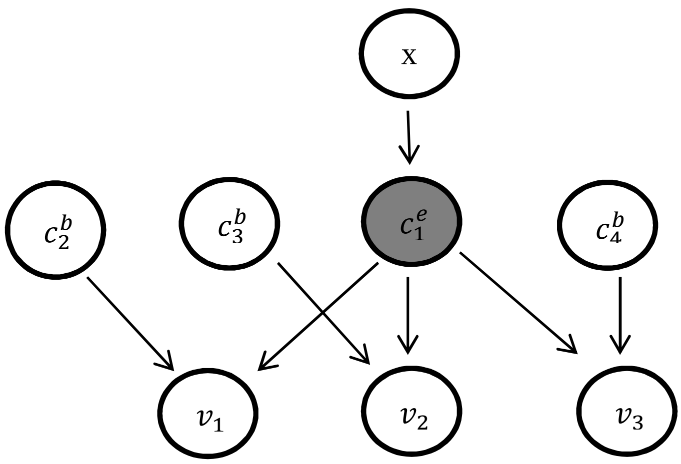

In figure 4.8 every agent has the same background factors, but different agents receive different pieces of evidence. Here majority voting would be an appropriate aggregation procedure. An example of this type of causal network might be where a group of climate scientists all have identical background education and training, but are able to make different experimental observations.

Figure 4.8: a causal network with common background factors, but private evidence.

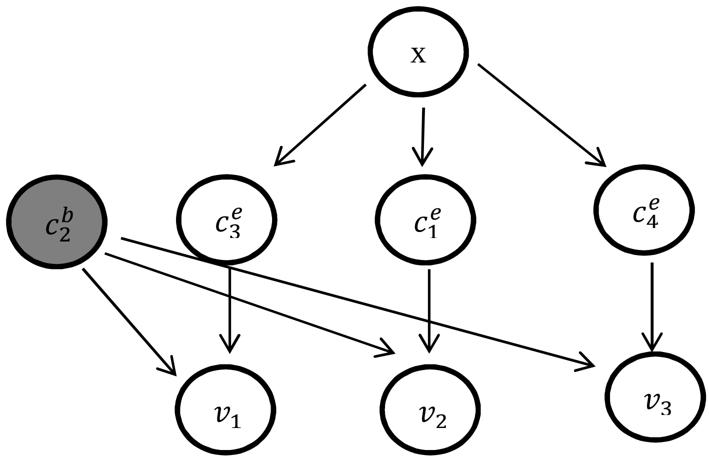

In figure 4.9, agent 3 receives all the information available. The only information 1 or 2 have is already possessed by agent 3. If these represent all the causes then the conditional probability of 1 or 2 voting correctly given the vote of 3 will be 1 or 0 i.e. there is no randomness in their probability of voting correctly, conditional on common factors. In this example, it may be more appropriate to make 3 the dictator rather than rely on majority voting involving all three agents. This is provided of course that agent 3 can prove to the others (of the social planner) that she knows something they don't, that she can prove her competence is the highest in the group.

Figure 4.9: a causal network where a single agent receives all the information.

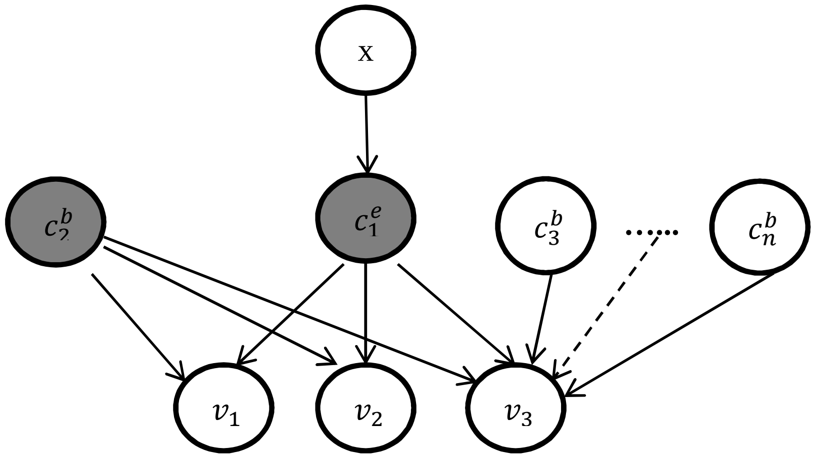

Finally, figure 4.10 will be more typical. Here every agent has four causes, two evidential causes and two background causes. Each agent has one evidential and one background cause in common, but because they have private evidential and background information, majority rule can aggregate the information.

Figure 4.10: a causal network with both private and common evidence and background factors.

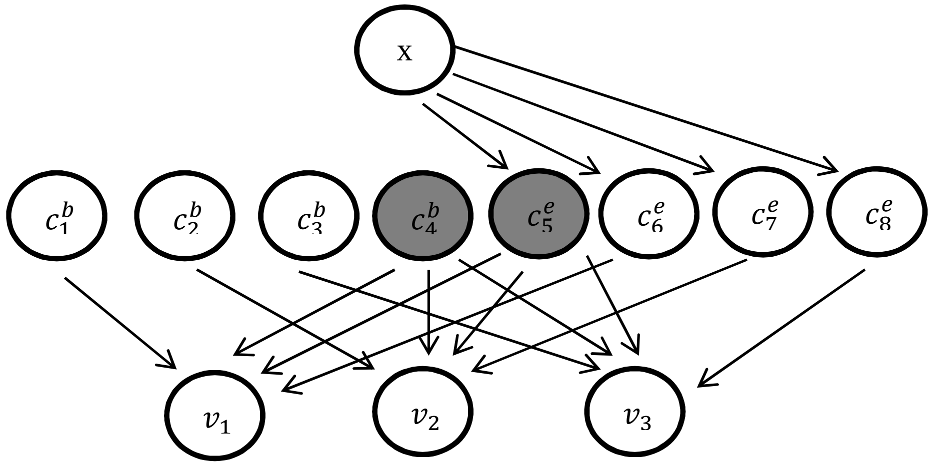

Summary

This chapter has focussed on how the inputs to judgement aggregation procedures form. It was argued that whether competence is transparent or not depends on contingencies in the circumstances of particular social choice problems. An agent will have competence $p_i > 1/2$ if they receive evidential signals and if the combination of evidential and background causal factors influencing their vote are overall truth-conducive. The votes of agents will be independent (and have some randomness), conditional on common factors, if agents have at least some causal factors influencing their votes (either evidential or background) that are held privately by them.

But agents do not have evidential and background factors a priori. To complete the justification of aggregation procedures as truth-tracking mechanisms we need to show how it is plausible that individual agents within a group will identify evidential and background factors and how they will identify factors not held by other agents. In chapter 5 I present a general model of a group search procedure by which agents can search for and identify objects. These objects could be the truth-conducive pieces of background and evidential information.

Chapter 5: Group search procedures.

This chapter provides a general framework for search procedures involving groups of agents. A single agent searching for an object of interest may only have a small probability of finding it. But if we employ a group to search for the object the probability that at least one of the group members will find it can be significantly higher. There are two different social epistemic mechanisms behind the epistemic performance of a group search procedure. Firstly, increasing the number of agents can increase the number of locations visited by the group. Secondly, increasing the number of agents can increase the probability that the object at a particular location will be recognised by a member of the group.

I present a theorem that states under certain assumptions the probability that a group of agents will identify a particular object is increasing in group size and in the limit tends to certainty. The assumptions of the theorem are modified to produce extensions of the theorem.

I then develop a model of a group search procedure to investigate the dynamics of group search. In the model there is a set of locations, one of which contains the object of interest. Individual agents engage in a search for the object by moving from location to location. The locations an agent visits are determined by four agent-specific variables: the agent's initial partitioning of the search space, the convention the agent employs for ordering the locations, the start *point of the agent's search and the agent's search heuristic. The objects an agent finds are determined by the locations the agent visits and the agent's capacity to recognise objects at those locations. If there are differences in the locations visited by agents and / or differences in the ability of agents to recognise objects then as group size increases the probability that a member of the group finds the object of interest increases and in the limit reaches certainty.*

The model of the group search is reproduced in the computer program 'NetLogo' and subjected to simulations. The results of the simulations confirm both the claims of the search theorem and the conceptual arguments of the search model: as group size increases the probability an object will be found increases and tends to certainty. The simulation results also show the impact on a group's search performance of adjusting the agent-specific search variables. I also present a proof of the theorem.

Search procedures

Suppose someone has lost their car keys. They may be able to narrow down where they lost their keys to their home. To find their keys a person needs to engage in a search, to move from location to location to see if that is where the keys were left. If their home is large it may take a very long time to find the car keys. And if we limit the time available for the search, for example if the person has to be at work in thirty minutes, then it is possible that they keys will not be found at all. But as we may know from our own experiences, the chance of finding a set of car keys can be improved if we increase the number of people looking for them. Someone who needs to find their car keys so that they can be at work in thirty minutes should ask the other members of their household to help search for them. Provided that each household member has at least some chance of finding the keys, and provided there are some differences in the way each agent searches, then a group of people will be far more likely to find the car keys than a single individual would be.

A search procedure does not have to be limited to the search for physical objects. In fact many of the more interesting applications of a search procedure apply to objects that are pieces of information. For example, suppose a philosopher is looking for an obscure Wittgenstein reference. They may look in the *Tractatus* and *Philosophical Investigations* but still not be able to locate it. Their chances of finding the reference will be improved if they email their colleagues asking for help. A group of philosophers are more likely than an individual philosopher to find the Wittgenstein reference, provided of course that each group member has at least some chance of finding it (for example each group member should be familiar with the main works of Wittgenstein) and provided there are some differences in the search behaviour of the philosophers.

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 5.1: search procedures.

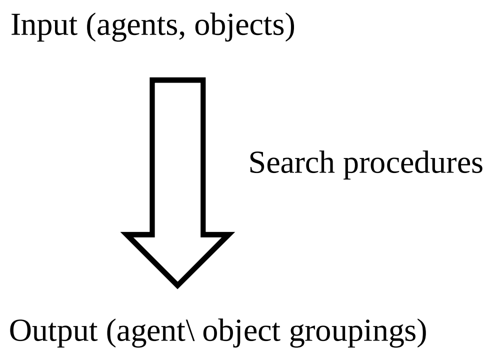

The epistemic advantage to including agents in the search for information has been investigated, for example, by Kitcher (1990, 1993) who provides a model under the title of a ‘division of cognitive labour’ to account for how different individuals in a scientific community spread themselves out over different possible avenues for research. Recently Weisberg and Muldoon (2009, 2011) have produced a model that more explicitly captures both the search involved in a division of cognitive labour, and the spatial aspect of a search. Weisberg and Muldoon have shown how it is epistemically desirable, from the groups’ perspective, to have a mixture of ‘maverick’ and ‘follower’ agents in the search for successful scientific approaches. 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 exploit the areas of epistemic significance identified by maverick agents. 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 will be identified.

The aim of this chapter is more modest than those canvased in the papers above. I am not attempting to provide an analysis of the optimal institutional arrangements for certain search problems (although this is a very interesting area for future research, and should be possible given the framework I set out later in this chapter). Instead the aim of this chapter is to make the case that in search problems, ceteris paribus, it is epistemically advantageous to include as many diverse agents in a group as possible. This result can subsequently be applied to solve problems in epistemic accounts of democracy.

Search Theorems

There are two distinct possible explanations for how search procedures allow groups of agents to find objects that may be missed by individual agents. These two explanations depend on how the differences in the search behaviour of agents are interpreted.

Suppose the group of philosophers is searching the *Tractatus* for the Wittgenstein reference. On one interpretation of the differences in search behaviour, one agent will search point 1, while a different agent will search point 2, a further agent will search point 3, and so on. The diversity in the *locations* searched by different agents in the group mean that a group of agents as a whole find objects missed by individual agents.

Suppose instead that each agent in the group of philosophers searches the entire contents of the *Tractatus*, from point 1 to point 7. One agent may read point 4.012 and fail to take in its significance. A second agent also reads point 4.012 but again does not recognise it as being important. A third agent reads point 4.012 and does recognise its significance. The differing abilities of agents to *recognise* an object at a particular location mean that a group of agents as a whole find objects missed by individual agents.

Similarly with the example of searching for car keys, there are two explanations as to why a group will be more likely to find the car keys than an individual agent will be. Firstly, there may be diversity in the locations searched by agents. If time is short and a single agent can only search part of the house for the car keys then as we increase the number of agents we increase the proportion of the house that is subjected to a search. Secondly, there may be differences in the ability of individual agents to recognise the car keys. Some agents may miss the car keys even if they visit the room where the keys are located. If all agents search the same locations in the house then, although some of the agents may fail to recognise the car keys this has no impact on the ability of subsequent agents to recognise the keys, and as the number of agents increases the probability that at least one of them will find the car keys also increases.

In what follows I present a series of theorems that capture in more formal terms the two kinds of search procedure. I begin by presenting the combined theorem before considering the spatial search theorem and search recognition theorem separately.

Combined Search Theorem

Suppose we have a set of objects O which are the subject of search. O may be known by agents (for example the car keys are known to exist); or O may be unknown (for example, agents searching for the Wittgenstein reference may not be aware of all the utterances of Wittgenstein). $o_1 \in O$ represents an individual object from the set of all relevant objects. We also have a set of agents $i, j, k, \dots, \in N$.

We have a finite set of locations L. $l_1 \in L$ represents a particular location from the set of all locations. Each location is atomic and cannot be divided into smaller locations. The set of locations L can be divided into jointly exhaustive subsets of locations[70]. Each of these subsets of locations are visited by different agents and are indexed according to the agents, namely $L_i, L_j, \dots \subseteq L$.

Each object occurs at a particular location. The mapping from the set of all objects to the set of locations is - initially - unknown to agents in the group. For example, l_{o_1} represents the location of object o_1 but the subset of locations $L_1, L_2, \dots \subseteq L$ in which o_1 occurs is unknown initially. We might think of l_{o_1} as

[70] The subsets of locations can, in some circumstances, contain only a single member (a single location).

a random variable (epistemically speaking) with a uniform distribution on L so that each $l \in L$ is equally likely to contain o_1.

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 $o \in O$ which is the subject of search. The location l^* of the object o and the subset of locations in which l^* 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 o. Once the agent moves to a location in their subset of locations, they attempt to identify the objects located there.

We write S_i for the event $o \in L_i$. We write $R_i|S_i$ for the event that an agent i recognises the object, given that the object occurs in their set of locations.

We will make four assumptions:

Spatial Search Competence: For each agent, the unconditional probability that the object occurs in the agent's subset of locations is uniformly bounded away from zero by some value α and is less than certainty. Formally, for all $i \in N$, $0 < \alpha < \Pr(S_i) < 1$.

Search Recognition Competence: The conditional probability that an agent recognises the object, given the object is in the set of locations visited by the agent, is uniformly bounded away from zero and less than certainty. Formally, for all $i \epsilon N, 0 < \gamma < \Pr(R_i | S_i) < 1$.

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

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

Combined Search Theorem

Given the above assumptions, the probability that a group of n agents finds the object is:

- *(non-limit claim)* increasing in group size; and
- *(limit claim)* in the limit approaches certainty[71].

The unconditional probability that a group of n agents finds the object o is given by[72]:

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

[71] This result is driven by the zero-one law, and not the law of large numbers as in the case of the CJT.

[72] Note that if the object does not occur in the subset of locations searched by an agent then there is no possibility that agent will recognise the object. More formally, $\Pr(R_i | \neg S_i) = 0$.

According to the combined search theorem, increasing group size is always epistemically advisable (provided the four assumptions hold). By increasing the group size we increase the probability that at least one of the group members will visit the object's location (they have at least some chance of recognising the object there). By increasing group size we also increase the number of agents visiting the particular location of the object and so increase the probability that the object will be recognised by at least one member of the group.

The Combined Search Theorem focuses on the search for a single particular object. We can apply the Combined Search Theorem to all objects $o \in O$. As such, as group size increases the total number of objects found should also increase.

To see the differences in the two types of search procedure, it is helpful to consider the two parts of the Combined Search Theorem separately.

Spatial Search Theorem

Here we abandon the Search Recognition Competence assumption from above. Instead we assume that the recognition competence of agents is perfect – if an agent visits the location of the object the agent is guaranteed to find the object[73]. We have two assumptions as follows:

- *Spatial Search Competence*
- *Spatial Search Independence*

[73] The Search Recognition Independence assumption now becomes irrelevant.

Spatial Search Theorem: Under the above assumptions the probability that a member of a group of $1, 2, \ldots, n$ agents finds an object of interest is:

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

If 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 initial presentation of the Spatial Search Theorem makes the general epistemic case for increasing the number of agents involved in a search. As we increase the size of the group we increase the probability that at least one of the group members will visit the location of the object (and recognise the object there). The Spatial Search Independence assumption ensures that if the object does not occur in the subset of locations searched by one agent, it might still occur in the subset of locations searched by a different agent.

If agents were to search exactly the same subset of locations, there would be no epistemic advantage to increasing group size, given the assumption that an agent will recognise the object if it occurs in the subset of locations they search. It is optimal for the group if each agent searches an exclusive set of locations. But if the sets of locations searched by agents are exclusive then the Spatial Search Independence assumption will not hold. For example, if one agent fails to find the object (because the object does not occur in that agent's subset of locations) it makes it more likely that a different agent will find the object. We can replace the Spatial Search Independence assumption with the following assumption, which captures the exclusivity of agent's searches.

Spatial Search Diversity: For any two agents $i, j: i \neq j$, the events S_i, S_j are mutually exclusive i.e. $S_i \cap S_j = \emptyset$. Informally, 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; 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: Informally, 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. $\Pr(\bigcup_{i=1}^{n+1} S_i) - \Pr(\bigcup_{i=1}^{n} S_i) > 0$.

Search Recognition Theorem

Here we abandon the Spatial Search Competence assumption from earlier. Instead we assume that the probability that the object occurs in the set of locations searched by each agent is certainty[74]. We have two assumptions as follows:

- *Search Recognition Competence*
- *Search Recognition Independence*

Search Recognition Theorem: 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; and
- (non-limit claim) in the limit tends to certainty.

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

[74] As such, the Spatial Search Independence assumption becomes irrelevant.

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.

If Search Recognition Independence is violated and we use Search Recognition Independence 2 then in calculations of the probability of a group recognising the object we must use the values for agents' search competence that are conditional on common factors.

<u>A taxonomy of searches, and trade-offs</u>

We can summarise the three types of search theorem in the diagram below:

Figure 5.2: a taxonomy of search mechanisms.

		Search Recognition	
		Yes	No
Spatial Search	Yes	Combined Search Theorem	Spatial Search Theorem
	No	Search Recognition Theorem	N/A

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

Both the Spatial Search and Search Recognition procedures could operate in political settings. Suppose that the House of Lords must consider whether banning a certain recreational drug will be effective at reducing the number of deaths that result from this drug. Suppose there is a crucial piece of information, which should inform the decision making: when the drug was made illegal in the US, organised crime took over production of the drug, the quality of the drug decreased, drug users were reluctant to seek treatment and as a result the number of drug-related deaths actually *increased.* It is unlikely that any member of the House of Lords would possess this information prior to the Bill being presented to them. We would expect that at least the cross-bench members of the House of Lords would undertake some research before casting their votes. The research can be construed as the conduct of a search procedure.

It may be that a number of the Members hear the same submission from a member of the public that cites the US evidence. There is every chance that a given member will fail to see the significance of the evidence – they may take a dislike to the member of the public presenting the evidence, or they may feel that the situation in the US is too dissimilar to the situation in the UK for the evidence to be relevant. Provided that there is diversity in the capacity of agents to *recognise* the evidence, then as the number of cross-benchers hearing the evidence increases, the probability that at least one of them will make use of the information in informing their judgement increases.

Alternatively, it may be that no member of the public proactively offers the US evidence to members of the House of Lords – the members have to search for the information themselves. Provided that there is diversity in the potential sources of information investigated by agents – for example if one Member consults the medical community, another consults the voluntary sector, while another consults the policing community - then as the number of cross-benchers searching for information increases, the probability that at least one of them will come across the US evidence increases.

When we apply the two search procedures to the example of the House of Lords and US evidence of criminalising a drug, we can see again that the two search procedures pull in opposite directions. If we encourage a number of Members to listen to the submission from the Commissioner of the London Metropolitan Police, they may not have time to consult representatives of the voluntary sector who may also have vital information. Similarly, if we encourage different Members to consult different sources, we decrease the probability that the Members will pick up all the relevant information from a given source.

A general model of a search procedure

Chapter 2 showed how groups of agents employing various aggregation procedures, such as expert dictatorship, negative reliability unanimity rule and majority rule can track the truth (identify the true state of the world) as group size increases. The Search Theorem also shows how groups of agents can accomplish epistemic tasks (find an object of interest) as group size increases. To investigate the normative implications of a search procedure we need to move beyond the Search Theorem and develop a more detailed model of the search. In particular we need a more detailed account of how it is that different agents are able to visit different locations.

The components of the model

Any model of real world phenomena has to trade off tractability against descriptive accuracy (including the accuracy of predictions). At one extreme, a model which captures all the features of the real world will be very accurate but very cumbersome to use. Similarly at the other end of a spectrum, a model that captures few features of the real world will be very easy to use but is unlikely to be accurate. The model of the search procedure I present here is intended to capture the variables that human agents would be aware of and would make use of. The choice of variables is not just for the sake of predictive accuracy, but also so that the model is a plausible representation of actual agent and group behaviour. Simpler models would be more elegant, but would not correspond as tightly to the target phenomena. The model of a search procedure presented here takes inspiration from the models presented in Hong and Page (2004) Weisberg and Muldoon (2009, 2011).

All objects of a search, be they concrete objects like car keys or informational objects like a reference from Wittgenstein, will occur at certain *locations*. For example the car keys might be located next to the telephone in the hallway; the Wittgenstein reference might be located at point 4.012 in the *Tractatus*. The set of all possible locations for an object of interest comprises the search space. The size of the search space, the number of objects in the search space, the particular locations of particular objects, and the recognisability of the particular objects at particular locations are all factors beyond the control of any agent involved in a search. The philosopher looking for the Wittgenstein quote has no control over where the quote is actually located, or how many of Wittgenstein's statements have been published, or how well the quote stands out. Similarly someone searching for the car keys has no control over the number of locations the car keys could occur at, or the clutter that obscures a view of the keys.

Although the search space will be beyond the control of agents, how they engage in a search is largely something that agents do have control over. The particular locations an agent visits are determined by four agent-specific search variables: an agent’s initial partition, locational convention, start point and heuristic.

From the perspective of an agent the search space could be too large, an agent may believe that some parts of the search space are more likely to yield the object than other parts, or an agent might think they need to double up on the search space visited by colleagues in case earlier searches missed objects. If an agent is going to engage in an effective search for an object they may choose to limit the number of locations they search by employing an *initial partitioning* over the set of all possible locations. For example, if an agent is searching for their car keys they may choose to only look in the rooms they visited since arriving home last night. Similarly, a philosopher may choose to limit their search for the Wittgenstein quote to the *Tractatus* since she does not have a copy of the *Philosophical Investigations* on her book shelf. In some cases, an agent may choose to search the entire search space. This may be because it is sufficiently small in size, or because they have no reason to believe any part of the search space is more likely to contain the object, or because they do not want to ignore any location that might contain the object.

If an agent is to engage in a methodical search of their partition they need to employ a *locational convention*, a way of ordering the locations in their partition. There is no objective locational convention, but some locational conventions will be of more use than others. For example, an agent who is searching for their car keys could divide up the surface area in their home into 10 cm^2 squares and order these squares according to a grid reference. A locational convention such as this could be communicated clearly to other agents and shared. Alternatively an agent could employ a locational convention based on the spatiotemporal locations on the path they took last night between when they locked their car and when they went to bed. Similarly, if an agent is looking for a Wittgenstein quote in the *Tractatus*, they could order the possible locations according to page number, or according to points 1 to 7. The ordering from points 1 to 7 is a more useful locational convention for groups since the page on which a quote occurs will vary according to the typesetting of a particular publisher.

The selection of a partition and imposition of a locational convention may occur simultaneously if an agent chooses a particular property to focus on. For example, if an agent chooses to focus on the property of being on the path they walked through the house last night then this simultaneously selects a subset of locations out of the search space and generates a spatiotemporal ordering over those locations. An agent could choose to search for the Wittgenstein reference by choosing the property of being in the *Philosophical Investigations*. In doing so they narrow down the number of locations they will search and they are presented with a ready-made ordering from page 1 onwards.

To commence a search within a partition an agent needs a *start point*. The partition and locational convention chosen by an agent might imply a certain start point. For example, if someone is looking for the Wittgenstein reference in the *Philosophical Investigations* then the natural place to start is on page 1. But many searches could begin at a random point on a partition. For example, if an agent is searching for car keys in a certain room, then any location in the room is an appropriate start point.

Once an agent has a partition, locational convention and start point they can begin searching for the object of interest by employing a certain search *heuristic*[75]. For example, suppose the agent is searching for the car keys, and they have decided to limit their search to their bedroom floor. The possible locations for the car keys on the bedroom floor have a natural two-dimensional ordering according to the width and length of the room. The agent has chosen the bedroom door as the start point. This agent might then employ a search heuristic of looking from left to right at every space on the floor as they walk forward from the door. Alternatively they could explore the edges of the room first, before moving inwards in a spiral.

Suppose instead that an agent is looking for the Wittgenstein reference. They have chosen an initial partition that limits their search to the contents of the

[75] The models of Hong and Page (2004) and Weisberg and Muldoon (2009) focus in particular on the specific types and combinations of heuristics that are optimal.

Philosophical Investigations, which is ordered according to page number, and they have decided to start their search on the first page. One possible heuristic is to search very thoroughly page by page, line by line. A different possible search heuristic is to read the first paragraph on each page and then read the remainder of the page that seems most likely to contain the reference.

The combinations of agent search variables

The combination of an agent's initial partition, locational convention, start point and heuristic determines the locations that he or she will visit. The objects an agent finds are determined by the locations he or she visits and the probability of recognising the objects at the locations.

If agents have same initial partition, locational convention, start point and heuristic then they will visit the same locations. It is also possible for agents to have different combinations of initial partitions, locational conventions, start points and/ or heuristics and yet still visit the same locations. For example, one agent might limit their search for the Wittgenstein reference to the *Tractatus*, order the content of the *Tractatus* according to points 1 to 7, start at point 1, and have a heuristic of looking at each point in turn. A different agent might also limit their search for the Wittgenstein reference to the *Tractatus* and order the content of the *Tractatus* according to points 1 to 7. This second agent chooses a start point of point 7 in the *Tractatus* and a heuristic of moving in a reverse numerical order. Although these two agents have different start points and different heuristics, at the end of the search process both of these two agents will have visited exactly the same locations. In this case the Spatial Search

Theorem does not apply. However, the Search Recognition Theorem may apply, provided that agent's search recognition competence levels are less than perfect (as per the Search Recognition Competence assumption) and provided that the recognition capacities of agents are independent (as per the Search Recognition Independence assumption).

It is also possible for the initial partition, locational convention, start point and heuristic of agents within a group to mean that some (but not all) of the locations visited by agents are identical. For example, two agents might limit their search for the Wittgenstein reference to the *Tractatus,* order the content of the *Tractatus* according to points 1 to 7 and choose to start at point 1. One of the agents chooses a heuristic of moving to the next prime number. This agent will then visit points {1,3,5,7}. The second agent chooses a heuristic of moving three points ahead. This second agent then visits locations {1,4,7}. In this simple example, the combination of agents' search variables means they visit some common locations {1,7}, but each agent visits some locations not visited by the other agent ({3,5} for the first agent; {4} for the second agent). If agents visit some locations visited by other agents, but also visit some of the locations not visited by other agents, then both the Spatial Search Theorem and Search Recognition Theorem of the Combined Search Theorem may apply. Here the probability of finding the objects is due both to the different locations visited by agents and to the differing ability of agents to recognise objects at the same locations.

Finally it is possible that agents have different initial partitions, locational conventions, start points and/ or heuristics such that they visit entirely different locations. For example, if we are coordinating a group of seven philosophers to search for the Wittgenstein reference, we might assign each philosopher a separate major point in the *Tractatus* to search. Here the Search Recognition Theorem does not hold but the Spatial Search Theorem may hold.

The set of locations visited by agents within a group would most likely only be identical or mutually exclusive if there is some coordination in the search behaviour of agents. In the absence of such coordination the intermediate case is most likely, where there is some overlap in locations visited by agents but each agent visits some unique locations. If a group of agents have some diversity in their initial partitions, locational conventions, start points and / or heuristics, they are likely to visit some different locations. As we increase the number of diverse agents we increase the number of locations visited by group members and therefore increase the probability that the object of interest will be found. In addition if agents recognition ability is independent then where there is overlap in the locations visited by agents we also increase the probability that the objects at the overlapping locations will be found.

Simulations

The general model of a group search procedure outlined in the section above was reproduced in the computer program NetLogo 4.1[76]. The assumptions of the search theorem were satisfied and the model was tested to see if it confirmed the theorems. Following the NetLogo conventions, the search space (set of all possible locations) or *initial partition* of the search space is represented in a two dimensional x and y tortoidal grid. The *locations*[77] in the grid are ordered (have a *locational convention*) according to a width and height coordinate. The grid is 37 locations wide and 37 locations tall meaning that there are a total of 1369 locations in the search space/ partition. Any of the locations could be a *start point* for an agent's search[78]. There are a variety of search *heuristics* agents could employ. For example an agent could rotate a random number of degrees to the right, and then move forward one location. Each simulation lasts 100 agent moves. In line with the simulations in both Hong and Page (2004) and Weisberg and Muldoon (2011) my simulations assumed that agents have identical initial partitions (to use my terminology).

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

[76] 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.

[77]Or 'patches' in NetLogo terminology.

[78] Agents are 'turtles' in NetLogo terminology.

Search Theorem where the Spatial Search Competence and Spatial Search Diversity assumptions hold. Thirdly, I present sample calculations for the Search Recognition Theorem. Finally I present simulation results for the Combined Search Theorem.

In each simulation model the number of agents in the group was varied, generally from 10 to 100 in intervals of 10 agents. The experimental result is the proportion of locations visited at the end of the 100 moves[79]. 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, before the simulation is run, is seen in the figure 5.3 below. Here ten agents are placed on the search space at random locations:

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

Figure 5.3: a screen shot of the first simulation, before agents move.

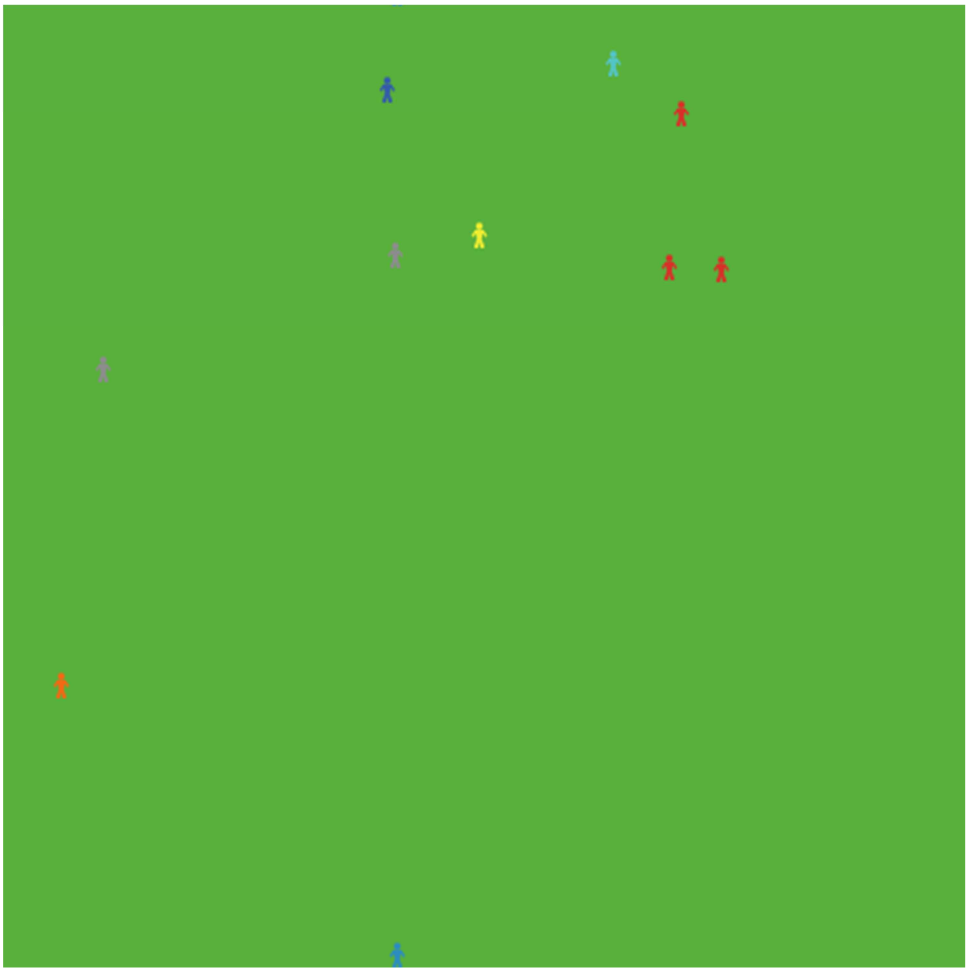

At the end of the simulation (at the end of 100 agent moves), the locations visited by agents have changed colour and the path taken by agents is traced. This is shown in figure 5.4 below:

Figure 5.4: a screen shot of the first simulation, after 100 agent moves.

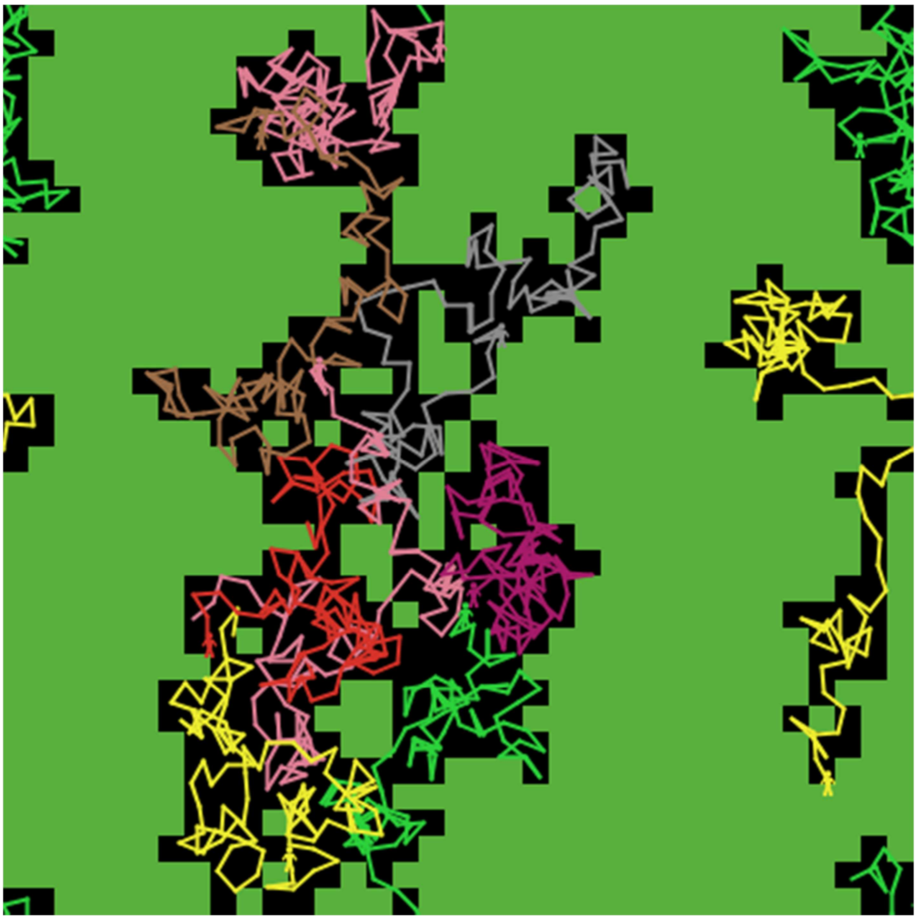

I now present the results of the two simulation models.

Simulation results

Simulation model 1: Spatial Search Theorem (Spatial Search Independence applies)

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. The start point of agents is 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. There is no restriction on agents exploring locations also visited by other agents, thus the Spatial Search Independence assumption holds. The results of the simulation are seen in the table and figure below:

Figure 5.5: simulation 1 results.

Group size	**10**	**20**	**30**	**40**	**50**	**60**	**70**	**80**	**90**	**100**
Proportion of locations visited	0.31	0.53	0.67	0.77	0.84	0.89	0.93	0.95	0.97	0.98

Figure 5.6: graph of simulation 1 results

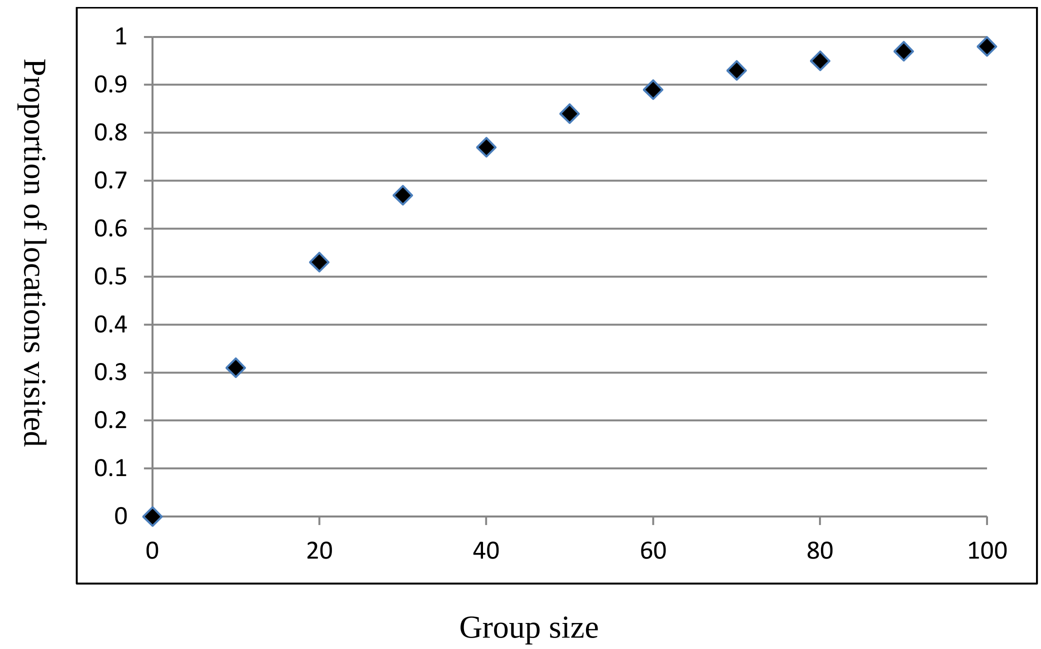

As can be seen in the figures above, the probability that an agent in the group will visit the location of the object (and by assumption find the object) is strictly increasing and in the limit tends towards certainty. This simulation provides confirmation for the Spatial Search Theorem.

Simulation model 2: Spatial Search Theorem, exclusive searches (Spatial Search Diversity assumption)

Here the start point of agents is determined randomly. Since each agent visits at least one location they have at least some probability of visiting the location of the object and so the Spatial Search Competence assumption holds. Each agent in the group employs the same type of search heuristic whereby they look at the location in front of them: if the location has not been visited before the agent moves forward; if the location has been visited before the agent rotates a random number of degrees to the right before looking at the next location in front of them. Thus the events of two agents visiting the same locations and identifying the same object are mutually exclusive and the Spatial Search Diversity assumption holds. The results of the simulation are seen in the table and figure below:

Figure 5.7: simulation 2 results.

Group size	**10**	**20**	**30**	**40**	**50**	**60**	**70**	**80**	**90**	**100**
Proportion of locations visited	0.15	0.27	0.35	0.41	0.47	0.51	0.53	0.57	0.59	0.63

Figure 5.8: graph of simulation 2 results.

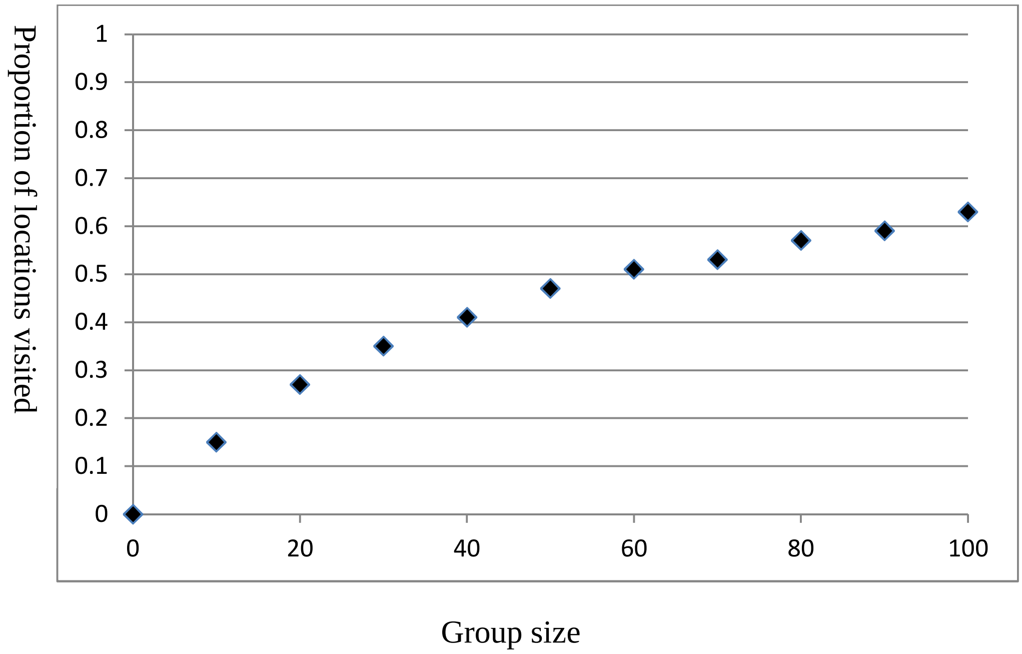

As can be seen, as group size increases the probability that each location is visited by at least one agent is strictly increasing and in the limit tends towards certainty. Thus the simulation results confirm the Spatial Search Theorem, this time where the search spaces of agents are exclusive and the probability of agents finding the object are not independent.

It is interesting to note that in this simulation the convergence towards certainty for finding the object was not linear but approximately exponential. This means that as group size increases the agents in the larger group do not have the same marginal capacity to search for objects that agents in smaller groups do. But this is to be expected, and is related to the finite nature of the search space. Suppose that all agents have an equal capacity to explore the search space. If we focus on the marginal contribution that each agent makes to the group search (the locations searched uniquely by the agents) then at most each agent explores

$\frac{1}{n}$ of the search space. As group size n increases, the proportion of the search space available to each agent, $\frac{1}{n}$, decreases.

Interestingly the convergence towards a certainty of finding the object is much more rapid in the first simulation than in the second. We can see why this might be the case by considering the screen shots of the second simulation in figure 5.9 below with the screen shot from the first simulation in figure 5.4.

Figure 5.9: a screen shot of simulation 2, after 100 agent moves.

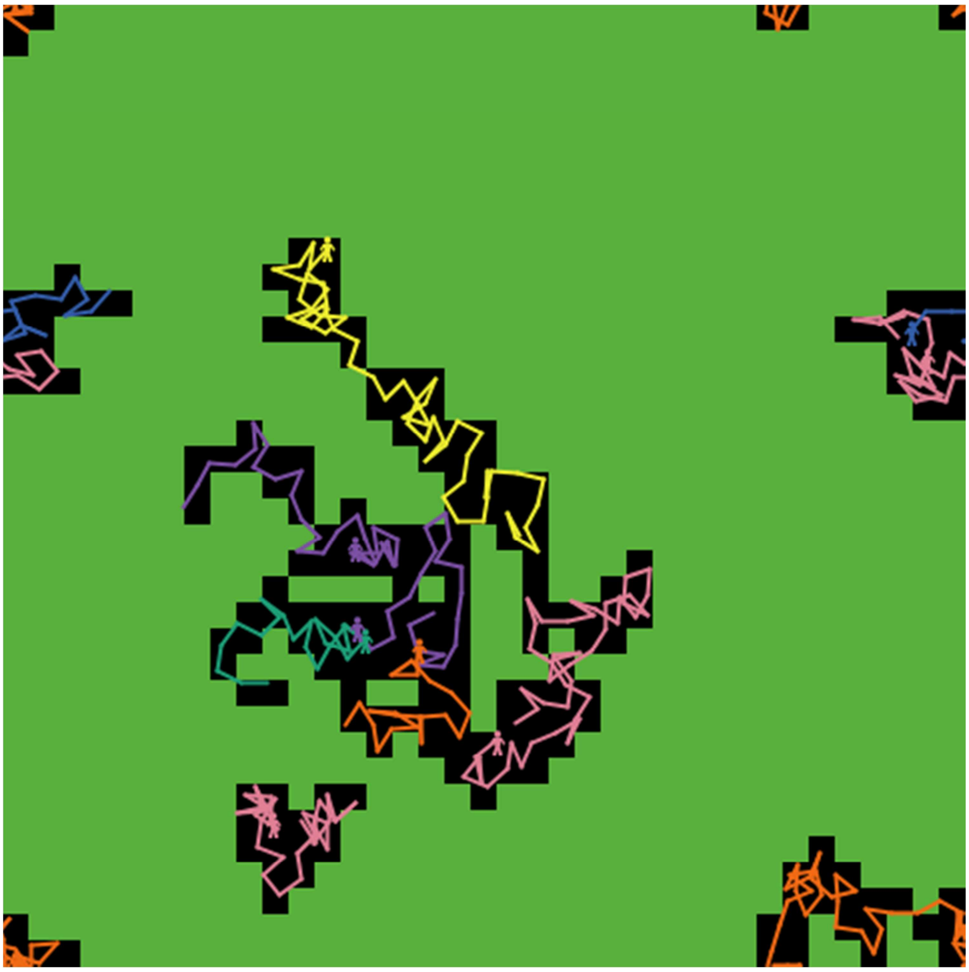

If agents are prevented from crossing into locations already visited by other agents (as per the second simulation) then this can box them in and limit the proportion of the search space that is accessible to the agents. We can see in figure 5.4 that agents are able to move quite some distance over the search space whereas in figure 5.9 agents are prevented from moving far when they run up against the search of other agents. In more concrete examples this shows that it can be desirable to allow some overlap in the locations visited by agents. For example, suppose two agents are looking for the Wittgenstein quote in the *Tractatus*. One agent starts at point 1 and employs a heuristic of moving to the location that is double their current point. A second agent starts at point 2 and employs a heuristic of moving 1, 2, 3 points ahead. If we permit agents to visit the same locations then the first agent visits locations {1, 2, 4} and the second agent visits locations {2, 3, 6} meaning between them they visit locations {1, 2, 3, 4, 6}. If however we prevent agents from visiting common locations then the first agent will be prevented from moving away from their start point and the total number of locations visited will be {1, 2, 3, 6}.

Similarly if a group of people are searching for the car keys they are more likely to be successful if they are not prevented from moving to other rooms to continue their search.

Sample calculations for the Search Recognition Theorem

The simulations presented thus far assess the Spatial Search Theorem in isolation. 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 in the calculations have reached the same location. Firstly we vary the value of search recognition competence $\Pr(R_i|S_i)$ to see the impact this has on the probability of a member of the group recognising the object. Secondly we vary the number of agents to see the impact that this has on the probability of a member of the group recognising the object.

Figure 5.10: the probability an object will be recognised, as recognition competence varies.

Note: group size fixed at 10 agents

$\mathbf{Pr}(\boldsymbol{R_i}|\boldsymbol{S_i})$	**0.01**	**0.05**	**0.1**	**0.2**	**0.3**	**0.4**	**0.5**
$\boldsymbol{P_R^n}$	0.096	0.4013	0.6573	0.8926	0.9718	0.9940	0.9990

Figure 5.11: the probability an object will be recognised, as group size varies.

Note: agent competence fixed at $\Pr(R_i|S_i) = 0.05$

$\boldsymbol{n}$	**10**	**20**	**30**	**40**	**50**	**60**	**70**	**80**	**90**	**100**
$\boldsymbol{P_R^n}$	0.401	0.642	0.785	0.872	0.923	0.954	0.972	0.984	0.990	0.994

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. For example, if an agent only has a 50% chance of recognising an object, if we place nine extra agents on that same location then it is close to certain that at least one of the agents will recognise the object at the location. Similarly when recognition competence is even lower and agents only have a 5% chance of recognising the object, when 50 agents visit that same location there is a better than 90% chance than at least one of the agents will recognise the object.

Simulation model 3: combined Search Theorem

Finally I present the results of a simulation that models the Combined Search Theorem. Here the start point of agents is 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. Firstly we set agents level of recognition competence to 0.5 (and so Search Recognition Competence holds[80]). The results are shown below in figures 5.12 and 5.13.

Figure 5.12: simulation 3 results, recognition competence $\Pr(R_i|S_i) = 0.5$.

$\boldsymbol{n}$	**10**	**20**	**30**	**40**	**50**	**60**	**70**	**80**	**90**	**100**
Proportion of all objects identified	0.22	0.40	0.53	0.63	0.71	0.78	0.83	0.86	0.90	0.92

[80] The probabilities of agents recognising the object are independent, so Search Recognition Independence holds.

Figure 5.13: a graph of simulation 3 results, recognition competence $\Pr(R_i|S_i) = 0.5$.

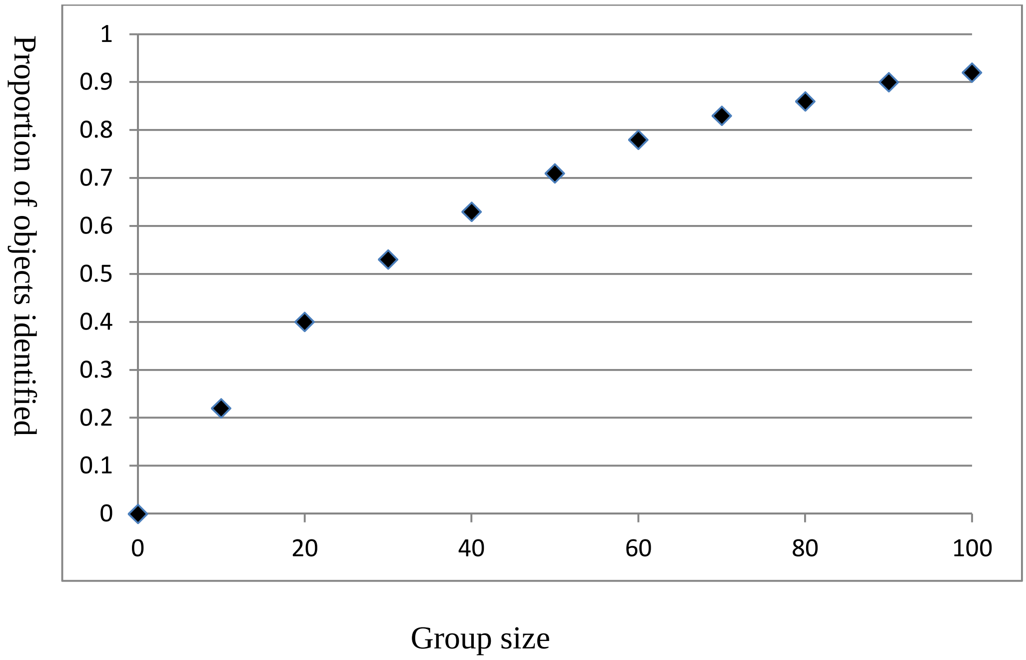

If we compare figure 5.13 with figure 5.8, we can see that (as expected), when the recognition competence of agents is less than perfect, the rate at which agents find the objects in a search space is much lower.

In the next simulation we set recognition competence to a much lower value of 0.05. Because the probability of agents recognising the object is much lower now group sizes need to be much greater if they are to be likely to find the object. Note that group size now ranges from 100 to 1000 agents. The results are seen in the two figures below.

Figure 5.14: simulation results 3, recognition competence $Pr(R_i|S_i) = 0.05$.

n	100	200	300	400	500	600	700	800	900	1000
Proportion of all objects identified	0.29	0.50	0.65	0.75	0.82	0.88	0.91	0.94	0.96	0.97

Figure 5.15: a graph of simulation 3 results, recognition competence $Pr(R_i|S_i) = 0.05$.

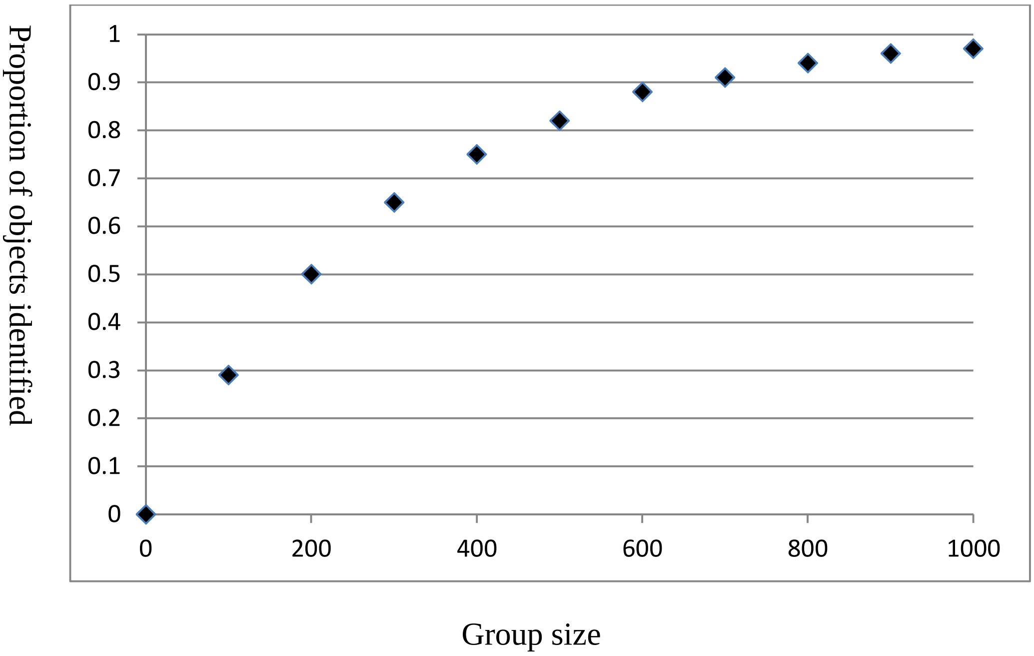

As can be seen, when recognition competence is low, it takes many more agents for the group as a whole to find the object. However these latter two simulations do confirm the more general Combined Search Theorem: the probability that a member of the group will identify an object of interest is increasing in group size and tends to certainty in the limit.

Proof of the Search Theorems

Here I set out a proof of the Combined Search Theorem. Please note that this proof is based on a proof by Franz Dietrich and so should not be considered original work.

We have four assumptions:

Spatial Search Competence (SSC): For each agent, the unconditional probability that the object occurs in the agent's subset of locations is uniformly bounded away from zero by some value α and is less than certainty. Formally, for all $i \in N,\ 0 < \alpha < \Pr(S_i) < 1$.

Search Recognition Competence (SRC): The conditional probability that an agent recognises the object, given the object is in the set of locations visited by the agent, is uniformly bounded away from zero and less than certainty. Formally, for all $i \epsilon N, 0 < \gamma < \Pr(R_i|S_i) < 1$.

Spatial Search Independence (SSI): The events of the object occurring in the subsets of locations visited by different agents are independent.

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

Combined Search Theorem

Given the above assumptions, the probability that a group of n agents finds the object is:

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

The unconditional probability that a group of n agents finds the object o is given by[81]:

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

We write F_i for the event that agent i finds the object o. An agent will find the object if and only if the object occurs in the subset of locations visited by the agent and the agent recognises the object.

The probability that agent i finds the object o is given by:

$$\Pr(F_i) = \Pr(S_i) \times \Pr(R_i|S_i)$$

Proof of the non-limit claim

For each n, $\bigcup_{i=1}^{n} F_i \subseteq \bigcup_{i=1}^{n+1} F_i$ hence by the monotonicity of probability, $\Pr(\bigcup_{i=1}^{n} F_i) \leq \Pr(\bigcup_{i=1}^{n+1} F_i)$.

[81] Note that if the object does not occur in the subset of locations searched by an agent then there is no possibility that agent will recognise the object. More formally, $\Pr(R_i|\neg S_i) = 0$.

Proof of the limit claim

The probability that n agents do not find the object is given by $(1 - \Pr(F_i))^n$. By SSC and SRC, $1 > \Pr(F_i) > \alpha\gamma > 0$. It follows that:

$$(1 - \Pr(F_i))^n \leq (1 - \alpha\gamma)^n$$

As $n \to \infty$, $(1 - \alpha\gamma)^n \to 0$. Therefore, as $n \to \infty$, $(1 - \Pr(F_i))^n \to 0$.

If the probability that n agents do not find the object tends to zero, the probability that at least one agent from a group of n agents do find the object tends to certainty.

Comments on search procedures

<u>Similar results from other models</u>

The Combined Search Theorem and model of a group search procedure show that there are epistemic gains from increasing the number of agents involved in the search for objects. As stated earlier in this chapter, the search theorems and model of a search procedure are inspired in particular by the models of Hong and Page (2004) and Weisberg and Muldoon (2009, 2011). These authors come to similar conclusions about the epistemic importance of increasing group size.

Hong and Page (2004) prove a lemma that a collection of agents will find the optimal solution (visit the location containing the object that is the optimal solution) with certainty as the group becomes large. If agents are drawn independently from a wider population it is unlikely that they will have common local optima (it is unlikely that they will have initial partitions, locational conventions, start points and heuristics such that they visit identical locations). Therefore as group size increases, the probability that they will have common local optima decreases to zero (the probability of all agents in the group visiting the same locations tends to zero).

Weisberg and Muldoon (2009) ran simulations that, inter alia, considered the epistemic impact of increasing group size. Firstly they considered ‘control’ agents who follow a search heuristic that pays no attention to the actions of other agents. In their simulations, a group of 10 control agents who have different randomly determined start points found the peaks on the epistemic landscape (found the locations with the objects or scientific discoveries of the most significance) 95% of the time. As group size in the simulations increased the probability that the peaks were discovered also increased, but with decreasing marginal returns. Furthermore when looking at areas of epistemic significance (at all locations with objects or scientific discoveries of importance, not just peaks) there is a linear relationship between the number of controls (agents) and average epistemic progress of the community (the percentage of locations with significant discoveries visited by agents). However progress at identifying these areas of epistemic significance can be slow since agents do not have the opportunity to learn from each other. These results again confirm the epistemic gains from increasing the number of agents involved in a search.

Interdependence of agents – balancing negative and positive correlations

Balancing the positive and negative correlations in the probability of agents visiting a particular location is important both in the model of the search procedure and the search theorem I presented. Group epistemic gain in the search for objects comes both from agents in the group visiting *different* locations and from agents in the group visiting the *same* location. When agents are guaranteed to recognise objects at particular locations then it is best if agents probabilities of visiting the same location are negatively correlated i.e. the fact that one agent visits a location should mean that a second agent will not visit that location. When it may be difficult for individual agents to recognise objects at particular locations there are epistemic advantages to positive correlations in the probability of agents visiting the same location i.e. the fact that one agent visits a location should mean that a second agent is more likely to visit that same location, since the more agents visit the location the more likely it is that at least one of them will recognise the object there. Of course positive correlation comes at the expense of negative correlation – we trade off the number of locations visited by the group as a whole against the increased likelihood of recognising the objects at the locations that are visited by the group.

The epistemic advantage of positive correlation in the search behaviour of agents is first identified in the model of List et al. (2008). List et al. provide an agent-based model of nest site selection by hives of honey bees. Their model proposes a mechanism to account for why the bee hives are so successful at identifying the best nest sites. In their model individual bee agents have a certain probability of flying to a particular nest site. The probability of finding the best site depends both on how likely the bee is to find the site on their own and the extent to which the bee finds the nest site based on the communication of other bees. Once the bees visit a potential nest site they assess the nest site's quality. A consensus for a particular nest site can emerge when more than twice the numbers of bees choose that site than the second most popular site and more than 20% of the bees choose that site.

The model of List et al. combines a search and aggregation procedure. The search procedure of their model includes both of the search mechanisms I have identified. There is a spatial component to search, where different bees may visit different locations or nest sites. There is also a recognition competent to the search. A single bee may be unreliable at recognising the quality of a nest site, but if lots of bees visit the same site there are lots of opportunities for the group to get an accurate reading of the nest site quality. The information the bees have regarding the quality of the nest site is then aggregated via the rules for consensus. The mechanism that allows the bees to pool their judgements regarding the nest site quality is Condorcetian in nature.

If the reliability of bee agents at recognising the quality of a nest site is better than random but not perfect (i.e. they have some ability) then the best way for the group to assess the quality of that nest site is by employing plurality rule. For plurality rule to track the truth the bee agents need to be *independent* in their assessment of the nest site quality. By assumption the bee agents are sufficiently competent (reliable) at assessing nest site quality. The remaining requirement for plurality rule to track the truth (to identify the true quality of the nest site) is increasing the number of bee agents visiting the site. For the number of bees visiting the site to increase, they need to be *interdependent*[82]: the fact that one bee visits the nest site and reports its location to the other bees should increase the probability that subsequent bees visit that same site.

The model of Weisberg and Muldoon (2009) also considers the balance between positive and negative correlations in the search behaviour of agents. While the search heuristics of control agents pay no attention to the behaviour of other agents, the heuristics of 'maverick' agents instruct the agents to avoid other agents (and so create negative correlations in the search behaviour of agents) and the heuristics of 'follower' agents direct the agents to move towards the searches of other agents (and so create positive correlations in the search behaviour of agents). The results of simulations show that homogeneous populations of followers are worse than homogeneous populations of control agents, who are worse in turn than homogeneous populations of mavericks at identifying the best scientific approaches (the objects of search). Homogeneous populations of mavericks, which are analogous to the agents in the Spatial Search Theorem, do quite well at identifying objects but not as well as heterogeneous populations of followers and mavericks[83]. It is best to have some

[82] List et al. acknowledge that total 'interdependence' (totally positive correlation) is epistemically bad for the group since it runs the risk of ignoring or missing the best sites.

[83] Note that Weisberg and Muldoon explicitly acknowledge the difference between the 'exploration' and 'exploitation' of a search space. Exploration involves moving from location to location (or in their model from approach to approach) whereas exploitation involves making full use of the scientific results at that approach (of recognising the object at that location). Their model focuses on the exploration of the landscape and assumes that the scientist agents will identify all the scientific results at the locations they visit. However there are still group epistemic advantages to positive correlations in the searches of the agents in the Weisberg and Muldoon model. This is because in the search space generated by Weisberg and Muldoon,

agents in the group who move away from other agents and find new areas of the search space *and* to have some agents who move towards other agents and look more carefully at the new areas of the search space.

The balance or trade-off between positive and negative correlations in the search behaviour of agents is a subtle one. Where the search recognition competence of agents is less than perfect, then a wholly negative correlation in the search behaviour of agents (so only one agent ever visits a given location) or a wholly positive correlation in the search behaviour of agents (so all agents visit the same locations) are both epistemically bad. In the former case, there is a real risk that the one agent visiting the location containing the object of interest will fail to recognise it. In the latter case, there is a real risk that the group will not visit the location of the object at all.

Negative and positive correlations in the search behaviour of agents may be more or less relevant at different times. For example, at the beginning of a search process agents will have no idea which locations are likely to contain objects. Therefore at the early stages of search there is no advantage to encouraging agents to visit the same locations (there is no advantage to positive correlations in the search behaviour of agents). In fact at the early stages of search there are some advantages to encouraging agents to spread out and visit as many different locations as possible (there are advantages to negative correlations in search behaviour). Once agents have engaged in search

locations with epistemically significant objects are clustered together. The positive correlations in agent search behaviour are not required so as to put multiple agents on the same location (as per the List et al. model) but rather so as to put agents on *adjacent* locations.

behaviour and have some idea which locations contain potential objects then there are advantages to encouraging fellow agents to visit those same locations (there are advantages to positive correlations).

There will be cases where a high degree of positive correlation in search behaviour is important. As both List et al. (2008) and Weisberg and Muldoon (2009) show, positive correlations in search behaviour can rapidly increase the *rate* at which objects are identified. This is because the positive correlations help concentrate the search behaviour of agents on promising locations (perhaps at the expense of missing some locations altogether). List and Vermeule (2010) show that for some types of problems this trade-off is desirable, in particular where there are time pressures on making a decision. For example a bee hive must choose a new nest site within a particular time period if it is to survive. Similarly, they cite the example of the US Supreme Court which considers on average 80 important cases each year. Arguably the processes by which cases are chosen by the US Supreme court closely resemble the nest site selection of a hive. If the selection of cases considered by the court was slower, if there was not the positive correlation in the court staffs' searches for suitable cases, far fewer important cases would be considered by the court. However List Vermeule (2010) suggest that in basic scientific research "…it is better that things be settled right, eventually, than things be settled today" (p.27). Positive correlation in the searches involved in basic scientific research would be epistemically harmful.

Communication between agents

Often encouraging positive correlations in the search behaviour of agents will require communication between agents. Bees, for example, are able to communicate the location and quality of the nest site via their dance behaviour[84]. Scientists are able to articulate their approaches (their research methods and results) in their journal publications[85]. However in some cases agents will be unable to clearly communicate the results of their search.

My model of search largely ignores communication between agents (although it is consistent with communication between agents). The reason that communication between agents is put to one side is that the model will be used subsequently to account for how agents find information such that the competence and independence assumptions of the CJT hold. The aim is to extend the explanation of how majority voting tracks the truth beyond the mechanism articulated in chapter 2 of this thesis. The CJT framework does not specifically require communication (or deliberation) and so the model of search which generates the judgement-generating factors of competence and independence for the aggregation procedures should not require deliberation (although deliberation is consistent with the search). The aim is to specify the minimal conditions under which a search will be successful, and an account of the mechanism by which the search will be successful.

[84] List et al. (2008).

[85] As per the assumptions of Weisberg and Muldoon (2009).

Summary

If the group task is to find a particular object, the two types of search procedure, as captured by the Spatial Search Theorem and Search Recognition Theorem respectively, are additive models of group productivity, according to the Steiner (1966) taxonomy. If one agent fails to find the object a different agents may be successful at finding the object. If we increase the size of the group then we increase the probability that one or other of the group members will find the object.

The next chapter will articulate how search procedures relate to judgement aggregation procedures.

Chapter 6: The link between search procedures and aggregation procedures.

Thus far this thesis has argued that truth-tracking by groups of political agents occurs via two procedures. Standard epistemic defences of democracy often focus on aggregation procedures such as majority rule, which pool the information individual agents have regarding the true state of the world. I also put the case for groups of agents employing search procedures to find information in the first place. The institutional features of search procedures and aggregation procedures amount to social epistemic mechanisms. This chapter discusses how the search and aggregation procedures link up. Search procedures allow groups of agents to extract information from the environment. But at the conclusion of a search the information will be dispersed across different agents. Aggregation procedures allow individual agents within a group to share the information they have extracted from the environment with the wider group.

The linking of search procedures to aggregation procedures fills two gaps in current epistemic defences of democracy that rely on aggregation procedures. Firstly, current accounts of aggregation procedures specify the types of judgement-generating factors (competence, independence, and transparency) required for a group to track the truth but they are silent on how the judgement-generating factors form. Search procedures can be used by agents to search for *evidential and background information to develop their levels of competence. Diversity in the search procedures of individual agents will generate the dependence relations in the group. Secondly, search procedures can be employed by a group to find possible alternatives and to set the agenda for a social choice.*

Once our framework for group truth-tracking joins search and aggregation procedures together we can consider the interaction between the two. We can see how contingencies in the way a search procedure is conducted mean particular aggregation procedures will be optimal at tracking the truth. Similarly, if an institutional decision is made in advance to use a particular aggregation procedure then this will influence the way in which a search procedure should be conducted so that it generates the appropriate levels and distributions of competencies and independence relations.

A two-staged process for group truth-tracking involving search and aggregation procedures

Chapter 2 of this thesis presented an analysis of different kinds of aggregation procedure, including expert dictatorship, negative reliability unanimity rule and majority rule, which generate a collective judgement or social choice as a function of individual judgements. Aggregation procedures have the ability to pool information contained in the judgements of individual group members.

However the aggregation procedures are silent on where the individual judgements come from.

Chapter 4 considered how the judgement-generating factors of competence levels, distribution of competencies, transparency of competence and independence relations form. It was argued that the probability that an agent votes for the correct alternative (the agent's competence) is determined by a combination of causes, including causes which are informational. These causes could be evidential or background, truth-conducive or misleading and private or common. The events of two agents voting for the correct alternative will be independent (conditional on the state of the world) if all their competence-generating causes are private. If agents have some competence-generating factors in common, independence can be regained by conditionalising on any common factors. An agent's level of competence will be transparent if the causal factors generating her competence can be shared with other agents or the social planner.

In the model I presented, the agents' final judgements are deterministic and are governed by the combination of causes they receive. The randomness in the judgements of agents, the reason we consider the probabilities of agents voting for the correct alternative to be strictly between 0 and 1, is that the observer or social planner is aware of some but not all of the causes of agents' judgements.

Chapter 5 set out a model for a group search procedure. The objects of interest occur at locations. The combination of an agent's search variables of initial partitioning, locational convention, start point and heuristic determine the locations the agent will visit. The locations an agent visits, along with their probability of recognising an object at a location, determine the objects he or she will find.

If the objects of search are the evidential and background information that produce the judgement-generating factors then by joining the search and aggregation procedures, we extend the explanation of how groups of agents can tack the truth. Information which could help an agent vote for the correct alternative is dispersed across a set of locations. To develop his or her competence an agent needs to engage in a search for information, to move from location to location to collect informational objects. The set of locations may be too large for a single agent to search on their own. Moreover, it may be that a single agent would have a probability less than certainty of extracting the information from the locations she visits. As such there may be epistemic gains to be had from increasing the number of agents involved in the search for information, provided of course that there are differences in the agents' search behaviour. Agents' different search variables mean different agents visit different locations and find different pieces of evidential and background information. Some overlap in the locations visited by agents mean the pieces of information missed by one agent may be picked up by another agent. The combination of information collected by agents from the environment produces the agents' judgement generating factors including their levels of competence. The information collected by agents can then be shared (indirectly) with the group via the aggregation procedure. The complete model of group truth-tracking is shown in the figure below.

Figure 6.1: the complete, two-staged, model of group truth-tracking.

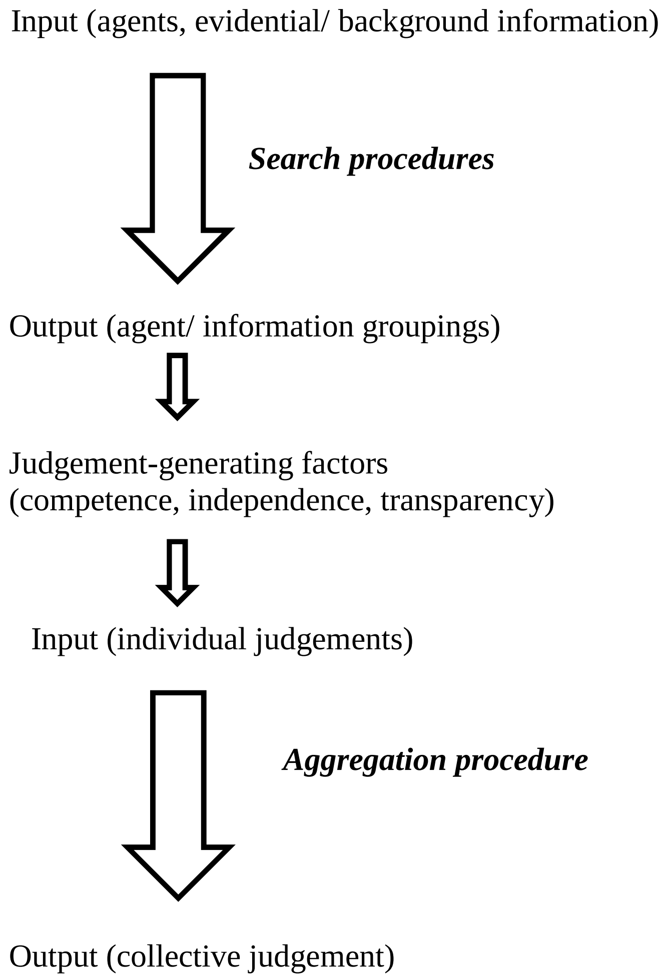

The competence of an agent will change over time as the agent moves from location to location and incorporates the information they receive at those locations[86]. The competence of agents at the time they cast their vote (express their judgement) will be determined by the combination of information they have received up till that point. The independence relations between agents will

[86] Or more particularly, if the competence of an agent represents the epistemic uncertainty of an observer or social planner then the competence of an agent will vary as the observer is aware of the agent moving from location to location and incorporating the information they receive.

also vary as they visit common locations and identify the same pieces of information. Again it is the independence relations at point in time which agents cast their votes which is of greatest interest.

Informally, the amount of diversity[87] in the group search procedure determines the independence relations in the aggregation procedure. This is because the amount of diversity determines the locations agents will jointly visit, which in turn determines the information held in common, which in turn determines the probability of voting correctly, given the votes of others. The amount of diversity in recognition capability also determines whether different agents visiting the same location find the same object.

Informally, the level of spatial search competence and search recognition competence determine the levels of aggregation competence[88]. This is because the level of search competence determines the amount of information an agent will find, which in turn determines the agent's level of aggregation competence.

A group's ability to track the truth is determined by the informational environment and the institutional features of both the search and aggregation

[87] The diversity in search behaviour was characterised by the following assumptions: Spatial Search Independence, Spatial Search Diversity, Spatial Search Diversity 2, Search Recognition Independence, and Search Recognition Independence 2.

[88] Here we assume that the level of spatial search competence and search recognition competence for finding a piece of information is identical for all pieces of information. It is more likely that an agent's ability to recognise a piece of information depends on what type of information it is. For example, some agents may be better at reading evidence than hearing evidence.

procedures. The quality and quantity of truth-conducive information[89] in the environment places a limit on the ability of groups of agents to track the truth. These limits are discussed in the next chapter. This thesis has already considered the institutional features of aggregation procedures and search procedures in isolation. This chapter will consider how the institutional features of search and aggregation procedures impact on each other. But before we consider this, we will pause to consider how the joining of search and aggregation procedures helps address two problems with the use of the Condorcet Jury Theorem (CJT) for epistemic defences of democracy.

Two problems with the CJT – setting the agenda and forming competence

The classic CJT provides a clear epistemic justification for widening the democratic franchise since as we include more and more voters in the group the probability of a correct majority verdict is monotonically increasing. But arguably the classic CJT rests on unstable foundations. Firstly the existing account of the CJT states that *if* the competence and independence assumptions hold *then* majority voting is likely to select the correct alternative as the social choice. In other worlds, the CJT only gives conditional support to the epistemic performance of majority voting. What is lacking is a justification for the antecedent of the conditional, a justification for how the competence and independence assumptions might plausibly be fulfilled. The current account begins at the point at which agents have already received private truth-conducive information such that the competence and independence assumptions

[89] Or more generally the amount of truth-conducive causal factors, including those that are non-informational in nature.

hold. But agents may not have private truth-conducive information, relevant to a particular agenda, a priori. We lack an account for how agents might come across pieces of truth-conducive information not shared by other agents.

The second concern with the foundations of the CJT is that the classic CJT is simply silent on how the two-placed agenda is set. The CJT implies that majority voting will tend to select the correct alternative as the social choice *if* the correct alternative is on the agenda. If the correct alternative is not on the agenda then the group will not be able to select it as the social choice. The epistemic challenge for a group of political agents is not just determining which alternative is correct when two alternatives are placed in front of them, but also determining which alternatives should be put in front of the group in the first place. Defences of epistemic democracy based on the CJT need an account for how the agenda is set in such a way as to include the correct alternative[90].

The model of a search procedure and the search theorem as presented in chapter 5 can be used to address these two problems with the CJT. In doing so we not only secure the foundations of existing epistemic defences of democracy based on the CJT, but deepen our explanation of how groups of agents are able to track the truth using majority rule.

The development of competence and independence in a CJT framework

Whether the competence and independence assumptions of the CJT hold is a contingent matter. In some social choice problems there will be misleading

[90] The problem of agenda setting is also raised by Fuerstein (2008).

evidence and so the competence assumption of the CJT will not hold. In other cases agents will have identical information and so independence assumption will be violated. If the competence and independence assumptions hold it will be because agents receive at least some truth-conducive information and they have at least some information not held by other agents. We could leave the analysis here, that it is just a happy accident that in a particular social choice problem the competence and independence assumptions happen to hold. But the group search procedure set out in the previous chapter can help explain *why* the competence and independence assumptions of the CJT hold, when they do hold.

Suppose there is a consensus in a parliament that the high level of drug use is causing harm to society. There are high levels of drug addiction, drug users are suffering health problems associated with use and there are high levels of crime attributed to drug use as it encourages the involvement of criminal gangs in supply and theft by addicts. Given a shared goal of reducing the harm to society caused by drug use, the members of a parliament need to implement policies that will be successful at reducing these costs. Suppose the most successful policy for reducing the social costs of drug use is decriminalisation and suppose the members of a parliament are facing an agenda of (decriminalise drug use/ do not decriminalise drug use). A vote will be taken to determine which alternative will be implemented and the vote will be decided by majority rule. The CJT implies that as more and more members are included in the voting group, the probability that the correct alternative of decriminalising drug use will be the majority winner increases. This is provided of course that the competence and independence assumptions of the CJT hold.

The competence assumption will hold in this example if the members of parliament are more likely to vote for ‘decriminalise drug use’ than for the other alternative. For this to be the case, each member must receive some truth-conducive information that tells them that ‘decriminalise drug use’ is the best alternative. This information could take the form of a fact finding mission to other countries that have decriminalised possession, or the advice of a local police officer who is concerned that arresting drug users makes the problem worse. Truth-conducive information could also include advice from medical professionals that drug addicts will not come forward for treatment if drug use remains illegal.

The independence assumption holds in this example if different members receive at least some different pieces of information[91]. For example, if one member receives the truth-conducive information from a fact finding mission to other countries, and a different member receives the advice of a local police officer then the votes of these members will be independent. If instead all members receive identical information, for example a submission from a medical expert, and this information determines their voting behaviour then we lose all randomness in the votes: the probability that one member votes correctly given that other agents vote correctly is certainty.

[91] Agents may have some common background factors (such as common beliefs) which means the events of agents voting correctly will not be independent conditional just on the state of the world. However, if agents have receive some different pieces of information then once the common factors are conditionalised on the agents will be independent and there will still be some randomness left in their votes.

For the competence and independence assumptions of the CJT to hold, the members of parliament need to search for truth-conducive information that will tell them which of the alternatives on the agenda is correct. Competence will hold if agents find some truth-conducive information. Independence will hold if agents find some information which is different to that found by other members. We can think of the pieces of information as having a location. The information that decriminalisation worked in other countries could be located within the government of that country itself, or in reports written about such case studies. The view of the local police officer on what works in reducing the impact of drug use can be accessed by visiting her at the police station, or via email.

The members of parliament can search for the truth-conducive information by choosing an initial partition, locational convention, start point and heuristic. If there is *diversity* in the initial partition, locational convention, start point and/ or heuristics that the members use then different agents will typically visit different locations and receive different pieces of truth-conducive information. If this is the case then the independence assumption of the CJT will hold when the members of parliament cast their votes. If the members of parliament have a certain level of *search competence* then as they search they will tend to find pieces of evidential and background information, their probability of voting for the correct alternative will be greater than a half, and so the competence assumption of the CJT will hold[92].

[92] In the model I present in chapter 4, the prior competence of agents is the probability that agents will vote for the correct alternative given no evidential information whatsoever. In a dichotomous choice the prior competence of agents will be $p_i^0 = 0.5$. If agents receive any piece of truth-conducive evidence as a result of their search their posterior competence will be

As the size of the group increases, the amount of truth-conducive information possessed by the group increases and in the limit all possible pieces of truth-conducive information are received by at least one member of the group. As such the group is able to extract the maximum amount of information regarding the true state of the world from the environment. This information can then be shared via the aggregation procedure of majority rule.

Agenda setting

The classic CJT begins at the point at which we have an agenda comprising two alternatives, one of which is correct. The classic CJT is silent on how this agenda is set. This should be of concern to anyone who wants to use the CJT to defend an epistemic conception of democracy. In the absence of an explanation for how the best alternative secures a place on the agenda, the CJT merely implies that the better of the two alternatives will be the social choice. And if the agenda is comprised of two mediocre alternatives, the majority winner will be a mediocre alternative.

As argued above if members of a parliament are faced with an agenda of (decriminalise drug use/ do not decriminalise drug use) and if the competence and independence assumptions of the CJT hold, then the parliament is likely to

greater than 0.5 and the CJT competence assumption holds. Of course it is not enough for the CJT that the competence of agents is greater than 0.5, we also need to establish that the distribution is symmetric about the mean. But if there are differences in the amount of information extracted by agents, as group size increases we should expect the distribution of competencies to begin to approximate a truncated normal distribution, which is symmetric about the mean. It is also possible that agents receive misleading information during the course of their searches, meaning their competence levels are below 0.5. The issue of misleading information is dealt with in the next chapter.

vote for the correct alternative: a majority of the parliament will vote to decriminalise drug use. But how do the alternatives of (decriminalise drug use/ do not decriminalise drug use) make it onto the agenda? Presumably if there is a consensus that *something* needs to be done to reduce the societal harm of drug use then there are any number of possible policy responses. Parliament could decide to increase the penalties for drug possession or increase the resources provided to police. More effort could be put on intercepting the supply of drugs into a country, or there could be more public advertising regarding the risks of drug use, or drug treatment services could be improved. The epistemic challenge is in fact two-fold. Firstly, the possible alternatives may not be immediately apparent to the members of parliament. They may need to do some work to find out the extent to which public advertising is a plausible solution to reducing drug use. Secondly, even when the members of a parliament are aware of all the possible alternatives, there remains a challenge of determining which two alternatives are the best and so deserve one of the two places on the agenda for a majority vote. We can think of agenda setting as a separate social choice problem in itself, with these two epistemic challenges of finding the alternatives and then choosing the alternatives to place on the agenda corresponding to a search procedure and aggregation procedure respectively.

A solution to this second epistemic challenge, of narrowing down the possibly large set of alternatives to the two allowed a place on the classic CJT agenda, comes from the various extensions of the CJT to cope with multiple alternatives, as discussed in chapter 3. List and Goodin (2001) extend the classic CJT to cope with multiple alternatives by moving from majority rule on a two-placed agenda to plurality rule on a many-placed agenda. Young (1998) details a Condorcetian extension of the classic CJT to multiple alternatives via a series of pair-wise social choices. In chapter 3 I argue that the List and Goodin extension risks decreasing the aggregation competence of agents and the Condorcetian extension can require too many elections. I argue for a mixed approach, with multiple elections, variable group sizes and agendas of multiple alternatives. If the concern is that we cannot narrow down our large set of possible alternatives to the two most likely to be correct, the solution is to place all possible alternatives on an agenda. This way we are guaranteed that if the correct alternative is identified by an agent, it will secure a place on the agenda. We can allow agents to place any plausible alternative they have found on the agenda because the costs of a placing the wrong alternative in the agenda are low, since once the agenda is set there will be a separate social choice to determine the final social choice alternative.

The remaining epistemic challenge is to find all the possible alternatives in the first place. This is where we can apply the model of the group search procedure. Although the set of possible alternatives are initially unknown to the members of the parliament, and as a consequence do not yet have a place on the agenda, each alternative will have a location. For example, the alternative of decriminalising drug use might have been put forward by a think tank in a 2009 report. The alternative of more police resources could be advocated by the chief of police. Public advertising may have been successful at reducing problem drinking in one city and so officials in that city would be able to propose this as an alternative for dealing with drug use. It is unlikely that a single member of parliament will be able to locate each of the alternatives, given the disparate locations of the alternatives. However the group as a whole may be able to identify all the possible alternatives by engaging in a search. The Search Theorem implies that as the size of the group involved in the search increases the probability that all the alternatives are identified is strictly increasing and in the limit reaches certainty, provided that each member of parliament has at least some chance of finding an alternative and provided that there is some diversity in the locations explored by agents.

The Spatial Search Competence assumption of the search theorem, as applied to this example, is fulfilled if each member of parliament engages in a search for alternatives. The crucial Spatial Search Independence assumption holds if there is some difference in the way members search for alternatives. If they all listen to the same submission at a hearing on drug use then the Spatial Search Independence assumption may not hold. But the Spatial Search Diversity 2 assumption of the Search Theorem would hold if the members of parliament choose different initial partitions. For example, if one member looks for legislative solutions and another member looks for best practice in other countries then, although there may be some overlap in the alternatives discovered by each agent, they are also likely to come up with some unique alternatives. The Spatial Search Diversity 2 assumption of the Spatial Search Theorem would also hold in this example if agents have the same initial partition but different start points, locational conventions and/or search heuristics. For example, the members of parliament may all share the same initial partitioning of parliamentary constituencies. But if each member begins searching for alternatives in their own constituency, then diversity is secured. Similarly diversity is secured if agents employ different heuristics in the search for alternatives. Perhaps one member searches for alternatives by holding an open meeting in their constituency, another member may search for alternatives by commissioning a literature review and a further member may search for alternatives by consulting public officials.

We could also apply the Search Recognition Theorem to the example of finding policy solutions for dealing with drug use. A select committee might hold hearings on the issue of the social costs of drug use and invite members of the community to offer solutions. Some members of the committee may ignore some of the proposed solutions if the solutions conflict with the member's political prejudices. However, if there is diversity in the political makeup of the committee, if different parts of the political spectrum are represented, then the different members will have independent recognition capacities and all proposed solutions will be recognised by at least one committee member.

Note that in applying the search procedure to the problem of identifying alternatives I have assumed that the objects of search are the alternatives themselves. However it could be argued that *all* objects of search, for any search procedure, are in fact evidential and background pieces of information. Consider the alternative of decriminalising drug use which is put forward by a think tank in a 2009 report. An agent who visits the location that is the 2009 think tank report will find the alternative of decriminalising drug use.

Equivalently, an agent who visits the location that is the 2009 think tank report receives evidential signals and background information which increases the probability that this agent will correctly judge that decriminalising drug use is an alternative that deserves a place on the agenda. Similarly, the alternative of more police resources could be advocated by the chief of police. If an agent visits this location (they communicate with the chief of police) then the agent will find the alternative of increasing the amount of police resources. Equivalently if an agent visits this location (they communicate with the chief of police) then they will receive evidential signals, auditory and visual pieces of information that are indirect causal relatives of the true state of the world, namely that increasing police resources is an alternative that deserves a place on the agenda.

Nested social choice problems

With the discussion thus far we have already seen that every social choice problem involves a two-staged process of a search procedure followed by an aggregation procedure. The search procedure allows agents to extract information from the environment. The aggregation procedure allows agents to share the information they have extracted with the group.

We have also seen that a given social choice problem may in fact exist within a nest, or sequence of other social choice problems. For example, the final social choice of a policy for dealing with the social costs of drug use can only occur after the separate social choice to determine which alternatives for dealing with the social costs of drug use should be placed on the agenda.

This nesting of social choice problems can in some cases be expanded beyond the two levels of setting the agenda and selecting the final alternative. For example, before a group addresses the problem of setting an agenda with policy alternatives for dealing with the social costs of drug use, it needs to come to a consensus regarding how the quality of policy alternatives are assessed. Some agents might believe that the success of a policy at dealing with social costs of drug use should be assessed solely in terms of the wider costs to the tax payer. A different agent might believe that the success of a policy at dealing with social costs of drug use should be assessed in terms of the impact on people's health. If different agents within a group assess alternatives according to different metrics we run the risk that the group will select the wrong alternative as the social choice[93]. The solution to this problem, of securing a common metric for assessing the quality of policy alternatives, is to have additional social choice exercises. Firstly, agents can search for possible metrics for assessing policies aimed at addressing the social cost of drug use, such as the wider costs to tax payers and the impact on people's health. These possible metrics can then be placed on a shared agenda. There can then be a separate social choice exercise to determine a common metric for the group. This separate social choice exercise will encompass both a search procedure whereby agents search for information which will tell them what metric on the agenda is best, and then an aggregation procedure where agents share their judgement as to which metric is best.

[93] For example if half of the agents use the metric of low costs to the taxpayer for measuring social costs then they will tend to vote for the alternative of public advertising. If the other half of agents use the metric of health costs for measuring the social costs of drug use then they will tend to vote for the alternative of increasing health resources. In these circumstances we would tend to get a tie.

The nesting of social choice problems can be extended even further. Before agents begin to consider which metric should be used to assess policy alternatives for dealing with the social costs of drug use, there needs to be agreement as to the policy area they are dealing with. A parliament only has a limited amount of time and so they must make a decision to focus on the issue of the social costs of drug use rather than, say, how to increase educational attainment or how to reduce the costs of health care, or the extent to which nuclear weapons should feature in defence policy. Again, this further social choice problem of determining the policy area that parliament will focus on can be addressed, firstly, by a search procedure for identifying possible policy issues and placing these on a shared agenda. Given a common agenda there can then be a separate search procedure where agents can search for information to tell them which of the policy areas on the agenda should be the focus of parliamentary time. The information obtained by individual agents regarding which policy areas are most important can be shared with the group via plurality rule: the policy area that receives the most votes in its favour will be the focus for the group.

In the remainder of this chapter we will consider the interaction of search and aggregation procedures: how institutional decisions over the choice of aggregation procedure impact on how search procedures should be conducted; and how the outputs of search procedures impact on the type of aggregation procedure that is optimal.

The relationship between search and aggregation procedures

<u>The choice of aggregation procedure and its impact on the group search procedure</u>

The institutional decision over which aggregation procedure to use may be made before a search procedure has been conducted and therefore before the type of judgement-generating factors of agents are formed. If the type of aggregation procedure has been set then we know the type of judgement-generating factors that are required for a group to track the truth (the level and distribution of competence in the group, independence relations and transparency of competence). This in turn implies what the search procedure needs to deliver in terms of the distribution of information across the agents in the group.

A group might make an institutional decision to use the aggregation procedure of negative reliability unanimity rule if there are high costs associated with wrongly deciding a false alternative is true. For example, it might be considered better to preserve a nation's status quo constitutional arrangements, even if they are not optimal, rather than take the risk of moving to new constitutional arrangements that are flawed.

A group might use the aggregation procedure of dictatorship if the costs of including a lot of agents in a decision are high and if it is likely that the dictator will have high competence. For example, a Government might delegate some decisions in the defence portfolio to the Minister of Defence. There are simply too many day to day decisions for the Cabinet to make as a group (via unanimity or majority rule). Moreover the Minister of Defence will have a large ministry to provide policy analysis to the Minister, so that her decisions are very likely to be correct.

Finally a group may choose to use the aggregation procedure of majority rule if there is symmetry in the costs of false positives and false negatives and if the costs of including a large number of agents in the group are sustainable. For example, for ordinary pieces of parliamentary legislation (such as liquor licensing laws) the costs of staying with the current bad legislation might be equal to the costs of moving to a new piece of bad legislation. Furthermore, requiring all members of parliament to vote does not impose additional costs as parliament has already set aside legislative time.

The combination of judgement-generating factors required for the different aggregation procedures to track the truth was discussed earlier in the thesis in Chapter 2, figure. 2.17

All three aggregation procedures considered in this thesis require the group search procedure to result in agents having at least some private evidential and background information. The institutional design of the search procedure should therefore simply aim for agents to find as much private truth-conducive information as possible.

It is perhaps best for a social planner (or a group of self-organising agents) to treat a search procedure as occurring in two phases. The distribution of information across the locations in a search space may be unknown ex ante. Similarly the extent to which information is easily extracted from locations may not be known ex ante. In the first phase of a search it may help to survey the entire search space to identify concentrations of information or locations where the information cannot be easily extracted by an agent. To optimise the initial phase of search the social planner should encourage some spatial search diversity in the group so as to identify the spread of information across the search space. The initial search should also involve some overlap in the locations visited by agents so as to identify the extent to which information is easily extracted from locations. Once the social planner has some idea of the distribution (and recognisability) of information in the search space they can organise subsequent searches so as to increase the amount of information extracted from the environment. For example, if some agents report that they find information at a particular location and other agents report that they found no information at this location then this suggests there is a need to organise subsequent searches so that there is some overlap in the locations visited by agents. Similarly if agents report that there is a cluster of locations containing significant amounts of information, the social planner may decide to concentrate the subsequent searches of agents on this patch of the search space.

What a search procedure can deliver to an aggregation procedure

It may be that a search procedure has already been conducted and that there is now a distribution of information across agents in the group. If the aggregation procedure employed by the group is sensitive to the distribution of information across agents in the group then the probability the group makes the correct social choice can be optimised. The same post-search distribution of information across agents could be caused by many different combinations of initial distributions of information in the environment and different mixtures of agent-specific variables of search (the initial partition, locational convention, start point and heuristic of agents). Therefore the following discussion of how institutional decisions regarding search procedures impact on aggregation procedures does not focus on the 'inputs' to the search procedure (the informational environment and the combinations of agents' search variables) but rather on the output from the search procedure.

The taxonomy of the distribution of information across agents set out below ignores whether the information is evidential or background, truth-conducive or misleading and common or private. Both evidential and background information can increase the competence of agents, as argued in chapter 4 of this thesis. The issue of misleading information is discussed in the next chapter. Whether information is common or private determines the independence relations between agents. I will argue in the next chapter that agent's level of competence should be conditional on the causal factors they receive. Thus, provided that we conditionalise on common factors, the violations of independence do not pose problems.

The outputs from the search procedure can be categorised according to whether agents find lots of information during their search, whether all or only some agents find information, and whether agents find partial or complete pieces of information. This is summarised in figure 6.2 below.

Figure 6.2: a taxonomy of the post-search spread of information across agents in a group.

Complete info / **Incomplete info**	**All agents have information**	**Some agents have information**
Lots of information	1 / 2	3 / 4
Little information	5 / 6	7 / 8

If agents find lots of information during the search their competence will be high; if agents find little information their competence will be relatively low. If all agents find information then every agent has something to contribute to the aggregation procedure; if only some agents receive information only some of the group members will have information to contribute to the aggregation procedure. Finally while many pieces of information will have truth-conducive value on their own, there may be some pieces of information that only have truth-conducive value in conjunction with other pieces of information.

To illustrate this final case, suppose that one agent discovers during their search for information that decriminalising drug use will decrease the cost of drugs. The fact that the cost of drugs will drop does not on its own imply that decriminalising drug use will decrease the social costs of drug use. Suppose a separate agent discovers during their search that most of the social costs of drug use are caused by addicts committing acts of theft to pay for their drugs. On its own this piece of information does not imply that decriminalising drug use will reduce the social costs. However in combination the two pieces of information discovered by two different agents do imply that decriminalising drug use will reduce the social costs.

If all agents receive some complete information (cell 5) then majority voting will be the optimal aggregation procedure. Here the information that agents receive means they have a competence level that is better than random but less than certainty. No single agent can reliably identify the true state of the world on their own. However by sharing their judgements with the wider group the aggregation procedure of majority rule can weed out the error in agent's judgements and utilise the information agents have found.

If some of the agents have a lot of complete information (cell 3) then dictatorship would be the obvious aggregation procedure (provided of course that their competence was transparent)[94]. Here the large amount of information received by the minority of agents means their competence levels will be high.

[94] Alternatively, the group could employ form of oligarchy, where only those agents with high competence (those agents who have received a lot of information) cast votes.

If many agents have lots of complete information (cell 1) then either majority rule or dictatorship are appropriate aggregation procedures. If only a handful of agents receive a small amount of information (cell 7) then there may simply be insufficient information for any aggregation procedure to determine the true state of the world. I consider this issue in more detail in the next chapter.

In cells 2, 4 and 6 there is sufficient information dispersed among agents that indicates the true state of the world. However the way in which the signals are dispersed in the group means that no individual agent has sufficiently high competence. Agents need to share the partial pieces of information if they are to make use of it, perhaps via a deliberative procedure. For example if one agent has discovered that most of the social costs of drug use are caused by addicts committing acts of theft to pay for their drugs and another agent has discovered that decriminalising drug use will decrease the cost of drugs, then the agents need to share what they know if the information is to have any truth-conducive value. Aggregation procedures, as characterised in this thesis, are unable to aggregate the truth-conducive information that does not directly increase agent's competence levels. This is because the aggregation procedures do not directly pool the information held by agents. Rather, the truth-conducive information held by agents generates the judgements of agents, which are then aggregated into the social choice.

The taxonomy in figure 6.2 helps illustrate three different reasons why we might want to consult a group in making a decision. Suppose a group of hikers is walking through a forest and they see movement ahead. If all agents in the group clearly saw what caused the disturbance (cell 1 in figure 6.3) then there is no need to consult the wider group. However if it is not obvious to all the agents what caused the disturbance they might consult the wider group. There are three explanations as to why the group as a whole may be successful in this task. Firstly, it may be that each group member got an independent view of the object that was good but not perfect. We ask each person what they saw. If one after the other says something like 'I think I saw a bear'...'It looked a bit like a bear'...'Maybe a bear'... then we will have a certain amount of confidence that the thing ahead is indeed a bear (this corresponds to cell 5, where the information can be pooled via the aggregation procedure of majority rule). Secondly, it may be the case that only one of our group members saw the object. The view of most of our group may have been obscured by foliage. But if we ask enough of our group eventually we will come across a group member who got a clear view of the object and is able to confirm it was a bear (this corresponds to cell 3, where the optimal aggregation procedure will be expert dictatorship). Finally, it may be the case that each member of the group only saw a part of the object (cells 2,4,6). For example, one person saw that the object had brown fur, so it had to be an elk, wolf or bear. Another person saw that the object had sharp teeth, so it had to be either a bear or a wolf. Finally someone else saw that the object was tall, so it must be a bear. In this last case agents must discuss what they saw for the group to identify the animal. Standard aggregation procedures cannot aggregate the truth-conducive information possessed by agents. Agents must share their information via deliberation of they are to identify the true state of the world.

There are three types of outputs from search procedures which deliver what is required for an aggregation procedure to track the truth. Firstly, where many agents have some complete pieces of information we can employ majority rule. Secondly where some agents have many complete pieces of information (or where many agents have many complete pieces of information) we can employ expert dictatorship. However where the group search procedure has produced an output where agents only have incomplete information, standard aggregation procedures will not track the truth. Agents need to share their incomplete pieces of information before they cast their votes.

Summary

Thus far the chapters in this thesis have discussed aggregation procedures, search procedures and the interaction between the two. We have considered the impact of institutional decisions regarding the choice of the aggregation procedure and the conduct of search procedure. We have also considered the agent-specific variables that are relevant for the aggregation procedures[95]; and we have considered the agent-specific variables relevant for the search procedures[96]. The remaining issue is how contingencies in the informational environment – in particular the possibility of misleading information and finite information – impact on the ability of a group to identify the true state of the world.

[95] The competence levels, independence relations and transparency of competence.

[96] Including spatial search and search recognition competence levels, and the spatial search and search recognition independence relations.

Chapter 7: The limits of the informational environment.

This final substantive chapter focuses on contingencies in the informational environment which place restrictions on the absolute epistemic performance of aggregation procedures.

I begin by addressing the problem of the possibility of misleading information. A small but significant literature on this topic has developed quite recently, focussed on majority rule and the Condorcet Jury Theorem. The analysis in this literature shows that the mere possibility of misleading evidence and background factors means the asymptotic limit of the Condorcet Jury Theorem is not certainty, but some value less than certainty. This means that the absolute epistemic performance of majority rule may be too low: groups using majority rule as an aggregation procedure may not be very reliable at identifying the true state of the world. The possibility of misleading information affects the aggregation procedures of expert dictatorship and negative reliability unanimity rule in different ways. Here I apply existing results to my two-staged framework.

I also address the problem of finite information. So far as I am aware, this problem has not yet been addressed by other authors. In some social choice problems there may simply be insufficient information for a group to determine *the true state of the world, no matter what search or aggregation procedures the group employs. I consider what institutional responses a group might employ to maximise the probability of a correct social choice when the amount of information is limited. In particular I argue for the sharing of information between agents, via deliberation. I show how the sharing of information can boost the levels of competence in agents, while still preserving independence relations.*

Search procedures allow agents within a group to extract evidential and background information from the environment. Aggregation procedures allow a group of agents to pool the information they have found during their search procedures. Institutional decisions over the conduct of the search procedure and over the type of aggregation procedure will affect the probability that the group identifies the true state of the world. For example, if it is known that the pieces of information can easily be extracted from locations then a coordination of the search procedures of agents to ensure they are mutually exclusive is optimal. Similarly if it is known that there is variability in the capacity of agents to extract information from locations, then coordinating the agent's search procedures to ensure some overlap in the locations visited would be epistemically advisable. If it is known, after the search procedure, that many agents have found pieces of information then majority rule may be the most appropriate aggregation procedure to employ. However, if at the end of the search procedure it is known that one agent has found all the relevant pieces of information, then expert dictatorship may be the most appropriate aggregation procedure.

The extent to which a group (or social planner) can make institutional decisions over the aggregation procedure to employ depends on the extent to which the competence of agents is transparent and the extent to which the group is aware of the independence relations. Even when a group is aware of these features, the choice of aggregation procedure can only influence the group epistemic performance to a certain extent. The level of competence of agents places an upper limit on the epistemic performance of aggregation procedures. The level of competence of agents is in turn (partly) determined by the search procedures.

The extent to which a group can make institutional decisions regarding the conduct of search procedures depends on the extent to which the group (or a social planner) is aware of the spread of information in the environment and the extent to which the group is aware of the search skills of individual agents. Even when a group is aware of these features, and coordinates the searches of individual agents optimally, there can be limits placed on the ability of groups to extract information from the environment. These limits are generated by features of the informational environment itself, by the quality and quantity of information available.

Chapters 2 and 5 of the thesis discussed aggregation and search procedures respectively. Chapter 6 discussed the interaction between search and aggregation procedures, how institutional decisions over search procedures impact on aggregation procedures and vice versa. This chapter (chapter 7) discusses the final limit that is placed on the ability of a group to track the truth.

Firstly, I briefly summarise the recent literature on the possibility of misleading evidence in a Condorcet Jury Theorem (CJT) framework. I consider the implications for this analysis on other aggregation procedures. I then move on to consider the issue of finite evidence and the institutional decisions that might be made to maximise the probability of a correct social choice where the amount of truth-conducive information is limited.

The possibility of misleading evidence

Majority rule and the CJT

The possibility of misleading evidence is a standard concern with the CJT. For example, Grofman et al. (1983) note that *if* the competence of agents is below 0.5 (*if* agents have on average received misleading information) then the probability of a correct majority verdict is decreasing in group size and in the limit approaches 0. Dietrich and List (2004) provide a new model of a jury theorem where the possibility of misleading evidence is conditionalised on. Here agents in a group do not have private pieces of evidence; rather the group shares a common body of evidence intermediate between the agents and the state of the world. If a group of agents faces a common, intermediate body of evidence then this places an upper limit on the probability of a correct majority winner. As the size of the group increases, more and more background information is added to the group so the ability of the group to appropriately interpret the evidence increases. However the probability that the group will identify the true state of the world (the probability of a correct majority verdict) is limited to the probability that the common evidence is non-misleading.

Dietrich and Spiekermann (unpublished a) produce a more general model. It is not just the possibility of misleading evidence that poses a problem for the classic CJT but the possibility of any sort of common circumstance (including both background and evidential information) that poses a problem. The classic independence assumption of the CJT requires that the events of agents voting correctly are independent conditional on the state of the world. Dietrich and Spiekermann note that this requirement will hardly ever be met – agents are likely to have at least some evidential and background factors in common. To recover independence in the votes of agents they propose a new independence assumption, which requires that the votes of agents are independent conditional on the state of the world and any factors held in common. The combination of the state of the world and common factors they term the 'problem'. With a problem-specific notion of independence secured, the classic CJT competence assumption needs to be revised. With the new independence assumption, the scope of the CJT has shifted to a fixed problem (with a fixed state of the world and a fixed set of (common) causes on agent's votes). The competence assumption must also refer to this same fixed problem, and the competence of agents must also be conditional on these (common) causes. The new competence assumption of their model requires that the problem specific competence of agents (the competence of agents given the state of the world and the common causes) is more likely to be greater than 0.5 than less than 0.5. In other worlds, the new competence assumption requires that the combination of common causes is more likely to be truth-conducive than misleading.

If the new independence and new competence assumptions of the revised CJT model hold, then the new jury theorem states that as the group size increases the probability of a correct majority verdict increases and in the limit tends to the probability that the combination of common causes is not misleading. As there is at least some possibility of a combination of misleading information, the maximum probability of a correct majority verdict in the new jury theorem is strictly less than certainty. A world in which most problems are misleading is unstable (Dietrich, 2008) and so the asymptotic limit of the new jury theorem is greater than ½.

The models of Dietrich and List (2004) and Dietrich and Spiekermann (unpublished a) show that the possibility of misleading information in social choice problems places restrictions on the absolute epistemic performance[97] of groups in a CJT framework. These restrictions, the fact that a group using majority rule may be significantly less likely than certain to identify the true state of the world, is not something that can be overcome. If it were possible to identify a piece or combination of information as misleading it would not be misleading. As Dietrich (2008) notes, a piece of information is misleading if receiving this information means an agent is less likely to vote for the correct alternative. To know if a piece of information is misleading the social planner needs to know what the true state of the world is. But if the social planner

[97] According to the definition set out in chapter 2, absolute epistemic performance requires that an aggregation procedure is good at identifying the true state of the world.

knows the true state of the world, there is no need to consult the judgements of individual agents.

The problem of misleading information is a particular problem for the aggregation procedure of majority rule because of the important threshold level of agent competence. If the combination of causal factors on agent's votes mean the competence of agents drops even slightly below 0.5, then as the size of the group increases the probability of a correct majority verdict tends towards zero. Neither expert dictatorship not negative reliability unanimity rule faces this problem, as I show below.

Negative reliability unanimity rule

Agents will have a prior competence of $p_i^0 = 0.5$ (according to my model assumptions, as discussed in chapter 4). Even if no agent receives any information whatsoever the aggregation procedure of negative reliability unanimity rule is increasingly likely to avoid the incorrect social choice as group size increases. The probability of avoiding the incorrect alternative is given by:

$$P^{-} = 1 - (1 - p_i)^n$$

If $p_i^0 = 0.5$ then:

$$P^{-} = 1 - (1 - 0.5)^n = 1 - \frac{1}{2^n}$$

and as $n \to \infty, P^{-} \to 1$.

Agents may receive truth-conducive or misleading pieces of information. Misleading information may mean that the competence of an agent is $p_i^1 < 0.5$. However, provided the information is not so misleading that the competence of agents is $p_i^1 = 0$, then misleading information does not pose a problem for the aggregation procedure of negative reliability unanimity rule. Provided that agents are not totally incompetent (provided that p_i^1 is not zero) then as $n \to \infty, P^{-} \to 1$.

The possibility of misleading information does not pose a problem for the baseline[98], relative[99] or absolute epistemic performance of negative reliability unanimity as an aggregation procedure. The upper limit of group epistemic performance is still certainty. Misleading information merely reduces the rate at which adding group members increases the probability of avoiding the incorrect social choice.

Expert dictatorship

The aggregation procedure of expert dictatorship only tracks the truth if the competence of agents is transparent. Here we need to be careful about how we interpret transparency. If competence is transparent, then 'misleading' information which decreases an agent's competence to below 0.5 is not misleading at all. An agent whose competence is $p_i = 0.5 - 0.2 = 0.3$ tells us

[98] As defined in chapter 2, baseline epistemic performance requires that a group using an aggregation procedure is better than random at identifying the true state of the world.

[99] As defined in chapter 2, relative epistemic performance requires that a group using an aggregation procedure is better than an individual at identifying the true state of the world.

just as much about the true state of the world as an agent whose competence is $p_j = 0.5 + 0.2 = 0.7$. If a group (or social planner) makes agent j dictator and accepts their judgement as the social choice then the probability of a correct social choice is 0.7. If the group instead makes agent i the dictator and makes the *opposite* of agent $i's$ judgement the social choice then again the probability of a correct majority verdict is $P = 1 - 0.3 = 0.7$. In fact, misleading information is just as epistemically virtuous for the aggregation procedure of expert dictatorship as truth-conducive information. The most epistemically difficult case is one in which agents receive no vote-determining causal factors at all and their level of competence remains at $p_i^0 = 0.5$. The most epistemically helpful factors are the ones that push an agent's level of competence as far away from 0.5 (in either direction) as possible.

The possibility of misleading information (that is, information which decreases the competence levels of agents) does not pose a problem for the baseline or absolute epistemic performance of expert dictatorship as an aggregation procedure. The upper limit of group epistemic performance is still certainty.

The real problem for the aggregation procedure of expert dictatorship is that the level of agent competence may not be *transparent*. This issue was discussed in chapters 2 and 4 of this thesis.

The problem of misleading information is related to the problem of finite information. If agents have access to all the information about the true state of the world then they cannot be misled as to the true state of the world. For the problem of the possibility of misleading information to have any purchase, it must be the case that either agents have access to an infinite but restricted pool of information, or agents only have access to a finite pool of information. The next section in the chapter is concerned with cases where agents have access to a finite pool of information but where there is no possibility of that pool of information being misleading.

Finite information

Chapter 4 of this thesis analysed how the judgement-generating factors of competence, transparency of competence and independence relations form. It was argued that agents need both evidential and background information for their competence to be better than random, for $p_i > 0.5$. In the absence of any evidential information whatsoever (including the absence of even the detail of the agenda agents are facing) agents have no clue which of the two alternatives on an agenda is correct and they will be forced to cast a vote at random. In these circumstances the competence of an agent will be $p_i^0 = 0.5$. Similarly agents need at least some background information to be able to interpret the evidence they have. For example, if an agent is told that the agenda comprises the alternatives (elephant/ not elephant) but the agent does not possess the background information that tells them what an elephant is, then again they will be forced to choose an alternative at random and their competence will be $p_i^1 = 0.5$.

If an agent has at least some evidential information and some relevant background information then his or her competence can be greater than 0.5 (provided that the combined effect of these factors is truth-conducive). For the aggregation procedure of expert dictatorship to have baseline epistemic performance[100] at least one agent has to receive enough background and evidential information (whose combined effect is truth-conducive) such that their competence is greater than 0.5. For the aggregation procedure of expert dictatorship to have absolute epistemic performance[101] at least one agent has to receive enough background and evidential information (whose combined effect is truth-conducive) such that the agent's competence is very close to 1.0. In addition, for expert dictatorship, the information received by the dictator must be of a kind such that the competence of the expert dictator is transparent.

For the aggregation procedure of majority rule to have good baseline and relative epistemic performance[102], agent's must have at least some evidential information and some relevant background information (whose combined effect is truth-conducive), and at least some of this information must be held uniquely by them[103]. If a given agent lacks evidential or background information their competence will be 0.5, and they will make no marginal contribution to group epistemic performance when added to the group (and in fact will add 'noise' to

[100] Baseline epistemic performance requires that a group using a particular aggregation procedure is better than random at selecting the correct alternative.

[101] Absolute epistemic performance means a group using a particular aggregation procedure is good at selecting the correct alternative.

[102] Relative epistemic performance requires that a group using a particular aggregation procedure is better than an individual at selecting the correct alternative.

[103] Here we rely on the CJT to support the claims of absolute and relative epistemic performance.

the group). If all the information an agent has is also held by another agent in the group then all randomness in vote of the new agent will disappear, the conditional probability of the new agent voting for an alternative given the vote of another agent is 1. Here again such agents will make no marginal contribution to the group epistemic performance when added to the group. For the aggregation procedure of majority rule to have absolute epistemic performance, the probability of a correct majority verdict must be close to certainty. For the probability of a correct majority verdict to be close to 1.0, there either needs to be an approaching infinite number of agents who have received enough private background and evidential information such that their competence is greater than 0.5; or alternatively (irrespective of group size) the agents must have received enough background and evidential information such that the competence of agents is close to 1.0.

A lack of evidential or background information is not a problem for the epistemic performance of negative reliability unanimity. Even if the competence of agents is $p_i^0 = 0.5$ then as the number of agents increases the probability of a correct social choice tends to certainty (assuming the votes of agents are conditionally independent). Negative reliability unanimity rule only fails to track the truth when the votes of agents are sufficiently dependent, or when the competence of agents is 0. For the posterior competence of an agent to drop from $p_i^0 = 0.5$ to $p_i^t = 0$ the agent must receive enough background and evidential information whose combined effect is severely misleading.

Finite evidence

There will be many cases where the amount of truth-conducive evidence available to agents is limited. For example, in a jury trial only a small amount of evidence (factors that are indirect causal relatives of the state of the world) may reach the jurors. There will only be so much forensic evidence left at the crime scene and only a finite number of witnesses to the actual crime. Let C^e represent the total body of finite evidence available as to the true state of the world. The truth-conducive strength of signals from the body of evidence depends on the prior competence level of agents receiving those signals. For example, it takes a *lot* more information to increase an agent's competence level from $p_i^1 = 0.7$ to $p_i^2 = 0.9$ than it does to increase an agent's competence level from $p_i^0 = 0.5$ to $p_i^1 = 0.7$. Suppose that if one agent i were to receive all of C^e then the agent's competence increases from $p_i^0 = 0.5$ to $p_i^1 = 0.9$. If two agents j, k each receive half of C^e then the competencies of these agents would increase from $p_j^0, p_k^0 = 0.5$ to $p_j^1, p_k^1 = 0.75$. I will briefly explain how the competence level for an agent with a portion of the body of evidence is calculated. $C_i^e = C_j^e \cap C_k^e$ and x is the true state of the world. The likelihood ratio (LR) = $\frac{\Pr(C_i^e|v_i=x)}{\Pr(C_i^e|v_i=\neg x)} = \frac{\Pr(C_j^e|v_j=x)}{\Pr(C_j^e|v_j=\neg x)} \times \frac{\Pr(C_k^e|v_k=x)}{\Pr(C_k^e|v_k=\neg x)} = LR^2$ by independence. To update the odds (in line with Bayes' rule): $Odds(v_i = x|C_i^e) = Odds(x) \times \frac{\Pr(C_i^e|v_i=x)}{\Pr(C_i^e|v_i=\neg x)} = Odds(v_i = x) \times LR^2 = \frac{0.9}{1-0.9} = \frac{0.5}{1-0.5} \times LR^2$. Therefore, $Odds(v_i = x|C_i^e) = 3 = \frac{\Pr(v_i=x|C_i^e)}{1-\Pr(v_i=\neg x|C_i^e)}$. Therefore $\Pr(v_i = x|C_i^e) = \frac{3}{4} = 0.75$.

We want to see the maximum possible level of epistemic gain for the group given the finite body of evidence, so we will assume that all pieces of evidence are truth-conducive (none are misleading). We also assume that no agents have factors in common and therefore agents are independent, conditional on the state of the world.

C^e is divided into n non-overlapping parts of equal strength $C_1^e, C_2^e, \dots, C_n^e$. Each agent i receives C_i^e.[104] The competence of agents depends on the pieces of evidence that they receive out of C^e. The prior competence of agent i is $p_i^0 = 0.5$. Agent $i's$ posterior competence given C_i^e is:

$$p_i^1 = \Pr(v_i = x | x, C_i^e)$$

Competence decreases as n increases because the amount of information in C_i^e decreases. Furthermore, as $n \to \infty, p_i \to 0.5$ because the amount of information in C_i^e tends to zero as $n \to \infty$.

On standard interpretations of the CJT finite evidence may pose a problem. As the number of voters tends towards infinity (as per the asymptotic CJT), voter competence tends towards 0.5. So as the number of voters tends towards infinity competence levels decrease towards a point where the competence assumption of the CJT no longer holds. If the competence assumption does not hold, the asymptotic CJT does not hold and we no longer have justification for the absolute epistemic performance of majority voting as an aggregation

[104] As such, agents receiving the evidence will be conditionally independent.

procedure. Put another way, according to the asymptotic CJT a requirement for majority voting approaching the upper epistemic limit of certainty is an approaching infinite number of voters whose competence is greater than 0.5 on average and who are independent in a formal sense. For this to be the case there needs to be an approaching infinite amount of evidence. If evidence is finite the upper asymptotic limit will never be approached. In cases where there is finite evidential information the upper asymptotic limit will always be some value less than 1.0. Furthermore, as the amount of evidential information decreases, the upper asymptotic limit of the probability of a correct majority verdict tends to 0.5. The CJT is inapplicable in the plausible cases in which competence levels depend on group size, or where evidence becomes common as group size increases, undermining independence.

Finite evidence also poses a problem for the aggregation procedure of expert dictatorship. Even if all the evidence available is given to the expert dictator there may be so little evidence that the competence of the dictator (and hence the probability of a correct social choice) is very low.

Finite background information

If finite evidential information poses a problem for the asymptotic limit of the CJT we can look to background information as a way out. Recall that the competence of an agent is determined by the combination of the evidential and background factors the agent receives. Although the background factors are by definition not causal relatives of the state of the world and therefore do not directly indicate which alternative on an agenda is correct, the background factors do have a causal impact on the votes of agents. Because background factors have a causal impact on the votes of agents they can either be misleading (if receiving a given background factor makes an agent less likely to vote for the correct alternative), or truth-conducive (if receiving a given background factor makes an agent more likely to vote for the correct alternative). While the amount of evidential information may well be finite (and in many cases extremely limited), there is likely to be much more background information that the group of agents can utilise. The background information helps agents interpret and make use of the evidential information and so obtaining background information can increase the competence of agents. If different agents obtain (or already have) different pieces of background information they will be (conditionally) independent.

For example, suppose a group of jurors must decide if a defendant is guilty of an act of murder. Juror $i's$ degree in biochemistry was in no way caused by the state of the world (the act of murder committed by the defendant). Furthermore, the fact that a juror has a degree in biochemistry is not indicative of the true state of the world. Nevertheless juror $i's$ degree in biochemistry helps an agent interpret the evidential information placed before them, such as DNA evidence. The juror's degree in biochemistry, combined with the DNA evidence, means that the competence of this juror will be greater than 0.5. A separate juror j may have studied botany rather than biochemistry, and as such has slightly different background factors than the first agent. The background in botany of juror j allows them to accurately interpret the DNA evidence, but in a way that is slightly different to the first juror. The background factors of an education in botany, combined with the DNA evidence, means that the competence of the second juror is also greater than 0.5. The slightly different way in which the second juror interprets the evidence means that the judgements of the two jurors are conditionally independent, i.e. the fact that the jurors have evidence in common means the fact that one votes for the correct verdict makes it more likely that a second juror votes correctly, however once we conditionalise on the common evidence the vote of one juror tells us nothing about how the second juror will vote.

Even if the amount of evidential information in an environment is finite and quite limited, there may be a significant amount of background information at the group's disposal. As such even a jury trial that suffers from a lack of evidence can be very likely to reach the correct verdict, provided that the jurors bring with them sufficient background information of sufficient quality such that they are able to correctly interpret the limited evidential information.

However the amount of background information may still also be finite. Let C represent the total body of truth-conducive information available in an environment. C encompasses both evidential and background information and C is finite. C is divided into n non-overlapping parts of equal truth-conducive strength C_1 , C_2 , ..., C_n . Each agent i receives C_i. The competence of agents depends on the pieces of information that they receive out of C. The prior competence of agent i is $p_i^0 = 0.5$. Agent $i's$ competence given C_i is:

$$p_i^1 = \Pr(v_i = x|x, C_i)$$

Competence decreases as n increases because the amount of information in C_i decreases. Furthermore, as $n \to \infty, p_i \to 0.5$ because the amount of information in C_i tends to zero as $n \to \infty$. The finite amount of truth-conducive information (evidential and background) appears to place a limit on the absolute epistemic performance of majority rule.

Limited background information also places a limit on the absolute epistemic performance of expert dictatorship. The amount of relevant background information may be so low that the competence of the expert dictator may be very close to $p_i^1 = 0.5$.

Possible solutions to finite information

We will now consider what routes there may be through this further bottleneck[105] to absolute group epistemic performance generated by finite information. To help illustrate the solutions I will assume that there are two groups of agents. These agents have extracted all the information from the environment that there is. One of the groups has homogeneous levels of competence of 0.6, while the other group has heterogeneous levels of competence which are symmetric about the mean competence of 0.6. The two groups are presented below:

$$(p_1, p_2, p_3) = (0.6, 0.6, 0.6)$$

$$(p'_1, p'_2, p'_3) = (0.4, 0.6, 0.8)$$

[105] The first bottleneck to the absolute group epistemic performance of a group is generated by the possibility of misleading evidence, as first identified in Dietrich and List (2004).

The first approach I consider for coping with finite information is the choice of aggregation procedure. The lesson from this approach will be that judgement aggregation procedures, such as majority rule, do not directly pool information. Rather, information generates the competence of agents which in turn generates the judgements of agents. The competence of agents should be interpreted as an epistemic capability. Secondly, I consider disaggregating a social choice and employing a premise-based aggregation procedure. The lesson from this second approach is that competence is a modular epistemic capability which can be used by the same agent on different agendas without violating independence. Thirdly, I consider the redistribution of information. The lesson from this approach is simply that it is indeed feasible for a social planner to allocate information across agents in a group as he or she pleases. Finally I consider the sharing of information. This combines the lessons of the three previous approaches, namely that competence is an epistemic capability that can be modular across agents. Because competence is a modular epistemic capability it can be used repeatedly to increase the probability of a correct social choice, even when the information generating the competence is limited.

The selection of aggregation procedure

Where the amount of information is finite we might be able to maximise its impact on the probability of a correct social choice by selecting the appropriate aggregation procedure. This point was made in the previous chapter where it was argued that the institutional decision over the aggregation procedure to employ should be sensitive to the post-search distribution of competencies in the group. First let's consider the case where the competence of agents is homogeneous. The probability of a correct social choice, given the aggregation procedures of majority rule, unanimity rule[106] and dictatorship are as follows:

$$P_{Majority\ Rule} = 0.648$$

$$P_{Unanimity\ Rule} = 0.216$$

$$P_{Expert\ Dictatorship} = 0.6$$

Clearly, when the competence levels of agents are homogeneous majority rule *is* the optimal aggregation procedure. Dictatorship is epistemically superior to unanimity rule since the event of a single agent voting correctly is more likely to occur than the events of all three agents voting correctly. Majority rule is epistemically superior to both unanimity rule and expert dictatorship because of its tolerance for mistakes. The correct alternative will be the majority winner if only two out of the three agents vote correctly. By contrast, with unanimity rule or dictatorship, if any of the voters make a mistake and vote for the wrong alternative the correct alternative will not be the social choice.

However, things may be different if the competence levels of agents are heterogeneous. Where the competence levels in our group are heterogeneous the probability of a correct majority verdict given the aggregation procedures of majority rule, unanimity rule[107] and expert dictatorship are as follows:

[106] Here we are considering the positive reliability of unanimity rule, not negative reliability unanimity rule.
[107] Here again we are considering the positive reliability of unanimity rule, not negative reliability unanimity rule.

$$P'_{Majority\ Rule} = 0.656$$

$$P'_{Unanimity\ Rule} = 0.21$$

$$P'_{Expert\ Dictatorship} = 0.8$$

If we compare the results for heterogeneous competence immediately above with the previous results for homogeneous competence we can see that both majority rule and expert dictatorship do better with heterogeneous competencies while unanimity rule does worse. Expert dictatorship does better with heterogeneous rather than homogeneous competence since heterogeneous groups contain individuals with high competence. Majority voting does better with heterogeneous competencies for small group sizes, but these differences will wash out as the size of the group increases.

When the amount of truth-conducive information is finite, it may be better in some cases to use expert dictatorship rather than majority rule as the aggregation procedure. If expert dictatorship is to be employed as the aggregation procedure then the competence of agents must be transparent. But if the competence of agents is transparent then we can choose to give more emphasis to the judgements of high competence agents and less weight to the judgements of low competence agents, in line with the following weights[108]:

$$w_i \propto \log(\frac{p_i}{1 - p_i})$$

[108] Grofman, B., Owen, G. and Feld, S.L. (1983) Theorem XIII, as discussed in chapter 2.

If we apply these weights to the aggregation procedure of majority rule then we get weighted majority rule, which represents the maximum possible probability of a correct social choice given fixed levels of competence. With weighted majority rule, any agent whose level of competence is not 0.5 can make an epistemic contribution to the group. With weighted majority rule the probability of a correct social choice is:

$$P_{Weighted\ Majority\ Rule} = \sum_{S \subset N} \prod_{i \in S} p_i \prod_{i \notin S} (1 - p_i)$$

where the sum is taken over all subsets $S \subset N : \sum_{i \in S} w_i > \sum_{i \notin S} w_i$. [109]

The probability of a correct social choice given the aggregation procedure of weighted majority rule is:

$$P_{Weighted\ Majority\ Rule} = 0.8$$

Majority rule can be thought of allowing a group to extract the signal from the judgements of agents while filtering out the noise (List, 2008). In the case of the group with heterogeneous competencies $(p'_1, p'_2, p'_3) = (0.4, 0.6, 0.8)$ the noise of agent 1 is drowning out the signal from agent 3. Applying weights in proportion to the competencies of agents means agent 1's vote is given a weight of $w_1 = -0.40547$ whereas agents 2 and 3 are given weights of $w_2 = 0.40547$

[109] Bradley, R. and Thompson, C. (2012), adapted from Owen, G., Grofman, B. and Feld, S.L. (1989), as discussed in chapter 2.

and $w_3 = 1.3863$ respectively[110]. Applying weights means we have greater confidence in the epistemic ability of some of the agents.

Judgement aggregation procedures do not pool information directly. Rather, information (evidential/ background, private/ common and truth-conducive/ misleading causal factors) generates the competence of agents. The competence of agents represents the agents' epistemic capabilities; it is a measure of their ability to identify the correct alternative on the agenda. These epistemic capabilities of agents can be utilised more or less effectively by different aggregation procedures. Weighted majority rule represents the maximum possible probability of a correct social choice, given fixed levels of competence. In employing weighted majority rule we acknowledge that some agents have greater epistemic capabilities than others (since they have received more truth-conducive information). Of course, if we choose to employ weighted majority rule, we dispense with equality of participation, which is also a virtue of democratic decision making.

Epistemic gains from disaggregation

The epistemic advantages to disaggregating a social choice into a set of premises and then holding majority rule decisions on each premise has been addressed by Bovens and Rabinowicz (2006), List (2006) and List (2008). Here I rehearse the main results and apply them to our group of a fixed size, where the amount of information is also fixed. I will argue that the mechanism that

[110] Note that in this particular example expert dictatorship is identical to weighted majority rule since the weight given to agent k is so much greater than that given to any other agent. However in other cases this will not be true. For example, if we have a group with heterogeneous competences of $(p_i, p_j, p_k, p_l) = (0.6, 0.6, 0.6, 0.7)$ then under the aggregation procedure of weighted majority rule no agent will be dictator.

accounts for the increased epistemic capacity that comes from disaggregation is that the competence of agents is a modular capacity that can be applied to different agendas. This important observation will be applied in the later solution for dealing with finite evidence: the sharing of information.

Any proposition is logically equivalent to a conjunction of other propositions. For example the proposition:

R = the defendant is guilty of manslaughter.

May be equivalent to:

P = the cause of death was blood loss due to being stabbed.

Q = the defendant stabbed the victim.

$(P\&Q) \leftrightarrow R$ = the defendant is guilty of manslaughter if and only if they stabbed the victim and this stabbing caused the victim's death.

The social choice can be made either by voting on the 'conclusion' R or by voting on each of the 'premises' P, Q and $(P\&Q) \leftrightarrow R$ and accepting the conclusion R if and only if P, Q and $(P\&Q) \leftrightarrow R$ are accepted. As authors such as List (2006), List (2008) and Bovens and Rabinowicz (2006) have shown, there can be epistemic gains from using a premise-based approach.

Agents will not have the same level of competence on conclusions as they will on premises. Knowing that the defendant is guilty is equivalent to knowing that the cause of death was blood loss due to stabbing AND the defendant stabbed the victim AND these two facts are necessary and sufficient for the defendant being guilty. Therefore we should expect that the level of competence on the premises will be higher than on the conclusion. If the homogeneous level of competence on the conclusion R is $p = 0.6$ then the homogeneous level of competence on each of these premises should be $\sqrt[3]{p} = \sqrt[3]{0.6} = 0.84343$. The probability that three agents with competence $p = 0.84343$ choose the correct alterantive via majority rule on a premise is $P_{Premise} = 0.93413$. The probability that agents choose the correct conclusion R, given a premise-based approach, is the probability that they make the correct majority choice on all three premises. This is given by $P_{Premise-based\ procedure} = 0.93413^3 = 0.81512$. We can compare the conclusion-based and premise-based decision procedures:

$$P_{Conclusion-based\ procedure} = 0.648$$

$$P_{Premise-based\ procedure} = 0.81512$$

In the case where competence levels are homogeneous there are clear epistemic advantages to a premise-based procedure.

Where the heterogeneous competence levels of agents on the conclusion are $p'_1, p'_2, p'_3 = (0.4, 0.6, 0.8)$, the competence levels of the three agents on the premises will be $p'_1, p'_2, p'_3 = (0.73681, 0.84343, 0.92832)$. The probability that this group makes the correct majority choice on a premise will be $P_{Premise} = 0.93461$ and the probability that the group makes the correct choice on the conclusion via a premise-based procedure is $P_{Premise-based\ procedure} = 0.93461^3 = 0.81638.$ [111]

We can compare the conclusion-based and premise-based decision procedures for groups with heterogeneous competencies:

$$P_{Conclusion-based\ procedure} = 0.65$$
$$P_{Premise-based\ procedure} = 0.81638$$

In the case where competence levels are heterogeneous there are clear epistemic advantages to a premise-based procedure.

Disaggregating a social choice problem into premises and a conclusion and then using a premise-based decision procedure is epistemically superior to both majority rule and expert dictatorship. What can account for the epistemic gain that comes from disaggregation, given that the inputs (the number of agents, and their truth-conducive factors) are fixed? I offer three explanations. Firstly, as noted by List (2006), the competence level on the conclusion can decrease rapidly as the number of premises increase. If competence drops below 0.5 then the probability of a correct majority will be less than the probability a single agent makes the correct judgement.

[111] Weighted majority rule makes no improvement in this particular case since the competence levels of agents on the premises are so similar.

Secondly, let a represent the number of premises. The probability of a correct premise-based verdict is given by:

$$P_{Premise-based} = (\sum_{h>\frac{n}{2}}^{n} \binom{n}{h} p^h (1-p)^{n-h})^a$$

The probability of a correct conclusion-based verdict is given by:

$$P_{Conclusion-based} = \sum_{h>n/2}^{n} \binom{n}{h} p^{ah} (1-p^a)^{n-h}$$

If we keep the level of competence on a premise p fixed but increase the number of premises a then this will have a detrimental effect on the probability of a correct social choice whether we use a premise-based procedure or a conclusion based procedure. Increasing numbers of premises decreases the probability of a correct social choice via a premise-based procedure because the final social choice is the *product* of the decisions on each premise. Increasing numbers of premises decreases probability of a correct social choice via a conclusion-based procedure because the level of competence on the conclusion is the level of competence on a premise *to the power* of the number of premises. The probability of a correct premise-based procedure is *always* greater than the probability of a correct conclusion-based procedure. This is shown in the figure below, where the number of agents is held fixed at three agents and the competence of agents on a premise remains fixed at $p = 0.6$. The figure shows that as the number of premises increases, the probability of a correct premise-based procedure (top line) will always be greater than the probability of a correct conclusion-based procedure (bottom line).

Figure 7.1: the probability of a correct social choice, given a premise-based or a conclusion-based procedure, $p = 0.6$.

Probability of a correct social choice

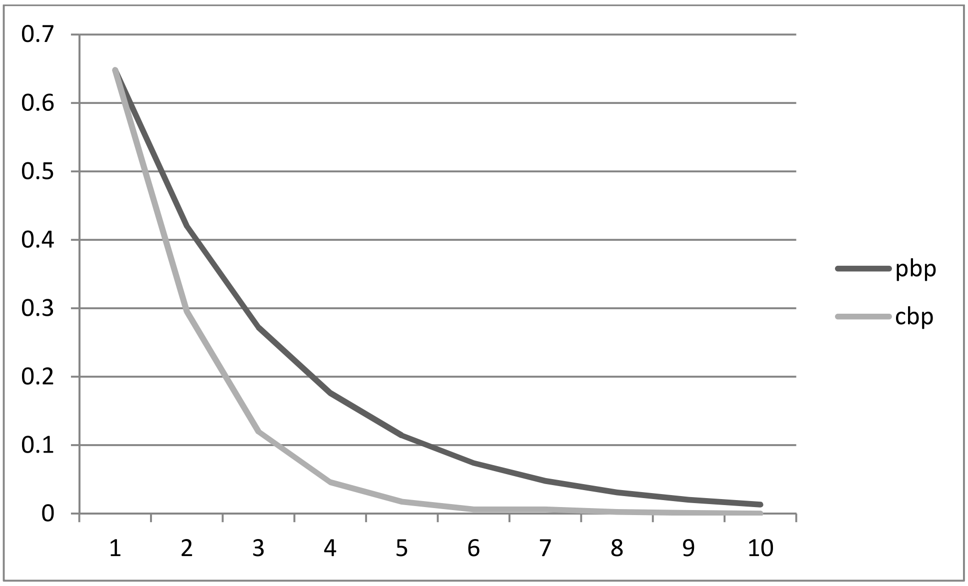

Number of premises

The third explanation for why a premise-based procedure is epistemically superior to a conclusion-based procedure is that competence can be modular. I have assumed in all the calculations above that votes of agents are independent across premises. For example, the fact that agent i votes for the correct alternative on the premise P (the cause of death was blood loss due to being stabbed) makes it neither more nor less likely that agent i votes correctly on premise Q (the defendant stabbed the victim). Bovens and Rabinowicz (2006) defend this type of proposition-wise independence on the basis of the modularity of competence. For example, an agent's physiological expertise at determining whether a victim died of blood loss (relevant for premise P) is different from that same agent's expertise at assessing witness statements that report the defendant stabbed the victim (relevant for premise Q).

I will argue that competence can be modular in a different sense. The same background factors can be used by the same agent on different premises to help them identify the correct alternative, while still retaining some premise-wise independence. For example, the background experience of paying close attention to complicated testimony may have a truth-conducive impact on agent competence levels on a variety of different agendas. Agent 1 might have a competence level of $p_1 = 0.84343$ on proposition P in part because of the background factor of experience of paying close attention to complicated testimony, in combination with the evidential factor of the testimony of the pathologist. Agent 1 might have a competence level of $p_1 = 0.84343$ on proposition Q, in part because of the background factor of experience of paying close attention to complicated testimony, in combination with the evidential factor of a witness statement. Because agent 1's votes on propositions P and Q have a common causal factor, the events of agent 1 voting correctly on propositions P and Q are not independent. The fact that agent 1 votes correctly on proposition P means agent 1 is more likely to vote correctly on proposition Q. However, the probability of agent 1 voting correctly on proposition P is independent of 1's vote on Q, conditional on the common background factor of experience of paying close attention to complicated testimony. Agent 1's vote on proposition P is in part generated by the evidential factor of the testimony of the pathologist. Agent 1's vote on proposition Q is in part generated by the witness statement. As such the vote of agent 1 on proposition P is independent of the vote of agent 1 on proposition Q.

This type of modularity of competence means that even when the amount of information is finite we can, in effect, increase the number of agents and increase the amount of information available to the group. In the calculations above there are three agents. With a conclusion-based procedure we aggregate one vote each from three agents. With the premise-based procedure, we aggregate three votes from three agents which (given conditional independence) is mathematically equivalent to aggregating a single vote from nine agents.

The law of large numbers, which accounts for the force of the CJT, is often explained by reference to coin tosses[112]. Suppose we have a slightly biased coin – perhaps there is a malfunction in the mint which creates a rounded edge of the coin on the 'tails' side, so that if the coin lands on its edge it will fall on 'tails' side not the 'heads' side. As a consequence of this bias the coin has a 0.51 probability of landing 'heads'. If the coin is tossed 100 times we should be very surprised if we get heads exactly 51 times. However if the same coin is tossed an infinite number of times we should get exactly 0.51 heads. The law of large numbers implies that the sample mean tends towards the population mean as the sample size increases.

We might use this slightly biased coin to decide if we should have an entrée at a restaurant ('heads' means 'yes' since we slightly prefer to have an entrée). We

[112] See for instance List and Goodin (2001) and Estlund (2008).

might also use that same coin to decide whether we should have a desert at the restaurant (again, 'heads' means 'yes' since we slightly prefer to have a desert). We are therefore 0.51 likely to have an entrée and 0.51 likely to have a desert. The fact that we order an entrée should make it neither more nor less likely that we order a desert, since each coin toss is independent.

The background factor of experience of paying close attention to complicated testimony is analogous to the defect in the coin. In each case the causal factor (the background experiences or the defect in the coin) generates a bias. In each case, the bias can have a causal impact on different agendas without violating independence across agendas.

Redistributing information

The next approach for increasing the upper limit of group epistemic performance, given finite information, is the redistribution of information. Thus far it has been assumed that our two groups of agents have already conducted a search procedure and all the information in the environment has been extracted by the agents in the respective groups. To recap, the distributions of competencies in the groups are as follows:

$$p_1, p_2, p_3 = (0.6, 0.6, 0.6)$$

$$p_1', p_2', p_3' = (0.4, 0.6, 0.8)$$

But agent competence levels do not have to be fixed. It may be possible for a social planner to control how the information is distributed across agents in a group. A social planner may be able to arrange a search so that only one agent receives all the information. Alternatively a social planner might organise a period of deliberation post-search so that the information extracted by agents can be redistributed across agents optimally.

A social planner might decide to maximise the competence of a single agent and make that agent the dictator. If a social planner can redistribute information across agents then there only needs to be enough finite information to increase the competence of a single agent from $p_i^0 = 0.5$ to $p_i^1 = 1.0$. If this agent i is made the expert dictator then the group will be able to identify the true state of the world with absolute certainty.

A social planner might choose instead to use the aggregation procedure of majority rule. While the CJT requires that the distribution of competencies in the group is symmetric about the mean, majority voting does not require any particular distribution of competencies. Grofman et al. (1983) Theorem IX shows the distribution of competencies that will maximise the probability of a correct majority where the amount of information is fixed:[113]:

a. if $p'n > \frac{n+1}{2}$ set a majority of individual agents' competencies to $p_i = 1.0$

[113] Grofman, B., Owen, G. and Feld, S.L. (1983) p.270.

b. if $\frac{n+1}{2} \geq p'n \geq \frac{n}{2} - 0.2$ set $p_i = 0$ for $\frac{n-1}{2}$ of the group and set $p_j = p'\left(\frac{2n}{n+1}\right)$ for the remaining $\frac{n+1}{2}$ group members[114]

c. if $\frac{n}{2} - 0.4 \geq p'n$, set $p_i = p'$ for all i

where p' represents the average competence levels and n represents the number of agents.

So the two groups of:

$$p_1, p_2, p_3 = (0.6, 0.6, 0.6)$$

$$p_1', p_2', p_3' = (0.4, 0.6, 0.8)$$

both come under the (b) category of Grofman et al. According to their approach we should adjust the competence of agents such that:

$$(p_1, p_2, p_3) = (0.0, 0.9, 0.9)$$

The probability of a correct majority verdict given this distribution of competencies is

$$P_{Majority\ Rule} = 0.81$$

[114] Note that there appears to be a typo in their proposal. They actually state $p_j = p\left(\frac{2n}{n+1}\right)$ i.e. they use a homogeneous level of competence. I presume they mean to use the average level of competence, which is more general.

However, I wish to note three problems with the Gofrman et al. approach. Firstly, there is a conceptual problem with adjusting the levels of competence as they propose. Note that in a dichotomous choice the prior competence of $p_i^0 = 0.5$ represents a position of ignorance, a situation in which the agent has no evidential information whatsoever. It will always be possible to shift an agent's level of competence back to $p_i^0 = 0.5$ by giving the information that would have gone to agent i to another agent. However it may not be possible to shift an agent's competence level to $p_i = 0.0$. An agent with competence $p_i = 0.0$ is entirely unreliable and just as valuable to a social planner as an agent whose competence is $p_j = 1.0$. It will only be possible to shift an agent's competence to $p_i = 0.0$ if agent i can be given the misleading information that would have gone to another agent. If there happens to be no misleading information in a group, no agent can have a competence less than 0.5 no matter how the information in the group is redistributed.

Under a more realistic interpretation of the Grofman et al. approach we can set the competence levels of a minority of agents to a position of ignorance $p_i = 0.5$ and redistribute the remaining fixed quantum of competence from these ignorant agents evenly among the remaining majority of agents. Following this approach means the competence of agents in our group will be adjusted to $p_1, p_2, p_3 = (0.5, 0.65, 0.65)$ and the probability of a correct majority verdict is:

$$P_{Majority\ Rule} = 0.65$$

It may be that when the point of ignorance is redefined as $p_i = 0.5$ rather than $p_i = 0.0$, that the calculations in the Grofman et al. approach also need to be revised.

The second problem with the Grofman et al. approach to a fixed sum total of competence (a finite amount of truth-conducive information) is that they presume additivity in the levels of competence. As I argued earlier, the truth-conducive strength of signals from the body of evidence depends on the prior competence level of agents receiving those signals. It takes a *lot* more information, for example, to increase an agent's competence level from $p_i^1 = 0.7$ to $p_i^2 = 0.9$ than it does to increase an agent's competence level from $p_i^0 = 0.5$ to $p_i^1 = 0.7$. The revised calculations in the Grofman et al. approach would need accommodate the fact that increased amounts of information have decreasing marginal impacts on an agent's level of competence.

The third, more significant, concern with the Grofman et al. approach to dealing with fixed amounts of competence is that it places an undue importance on preserving independence in the votes of agents, conditional just on the state of the world.

Sharing information

The section immediately above proposed redistributing information across agents in a group to maximise the possibility of a correct social choice. It was assumed that all information must be held privately by agents. For one agent to receive an extra piece of information it must be taken away from a separate agent. But there may be more significant gains to be had by agents sharing information. Sharing information would mean that all truth-conducive information would be held in common between agents.

Consider the case where agents have homogeneous levels of competence. In this case each agent received separate pieces of information that increased their levels of competence from $p_1^0, p_2^0, p_3^0 = (0.5, 0.5, 0.5)$ to $p_1^1, p_2^1, p_3^1 =$ (0.6, 0.6, 0.6). If all the information is given to a single agent then the competence of that agent will be $p_i^2 = 0.77$.[115] If this same information is also given to agent's j and k then the competence of agents j and k will also be $p_j^2 = 0.77$, $p_k^2 = 0.77$. If the social choice is to be decided by the aggregation procedure of majority rule then the probability of a correct social choice will be:

$$P_{Majority\ Rule} = 0.86563$$

The sharing of information, before judgements are aggregated via majority rule, is by far the most effective approach for maximising the probability of a correct social choice when truth-conducive information is finite. It makes more of a difference to the probability of a correct social choice than the choice of aggregation procedure, the use of a premise-based procedure or the redistribution of information.

[115] The likelihood ration (LR) = $\frac{\Pr(C_i^e|v_i=x)}{\Pr(C_i^e|v_i=\neg x)} = \frac{\Pr(C_j^e|v_j=x)}{\Pr(C_j^e|v_j=\neg x)} \times \frac{\Pr(C_k^e|v_k=x)}{\Pr(C_k^e|v_k=\neg x)} \times \frac{\Pr(C_l^e|v_l=x)}{\Pr(C_l^e|v_l=\neg x)} = \left(\frac{0.6}{1-0.6}\right)^3 = 3.375$. $Odds(v_i = x|C_i^e) = Odds(v_i = x) \times LR$. Therefore $\Pr(v_i = x|C_i^e) = 0.77$.

The calculations immediately above assume that the votes of the three agents are independent. There may be a concern that since agents have shared information, and now have information in common, the votes of agents are no longer independent. The more general worry is that if agents share all their causal factors they will have identical vote-determining causal factors and the votes of agents will be entirely dependent: there will be no randomness in the vote of an agent, conditional on the vote of another agent. As such, the probability of a correct majority verdict will be identical to the probability of a single agent voting correctly (the agent's competence) and adding agents to the group will make no difference. However I will show in the section below why the competence of agents, conditional on common factors, retains some randomness.

Case 1

Consider three agents i, j, k. The prior levels of competence for these three agents will be:

$$p_i^0 = \Pr(v_i = x|x) = 0.5$$

$$p_j^0 = \Pr(v_j = x|x) = 0.5$$

$$p_k^0 = \Pr(v_k = x|x) = 0.5$$

Suppose that these three agents receive similar (but different) packages of information, as represented in the figure below.

Figure 7.2: a causal network with three private packages of information.

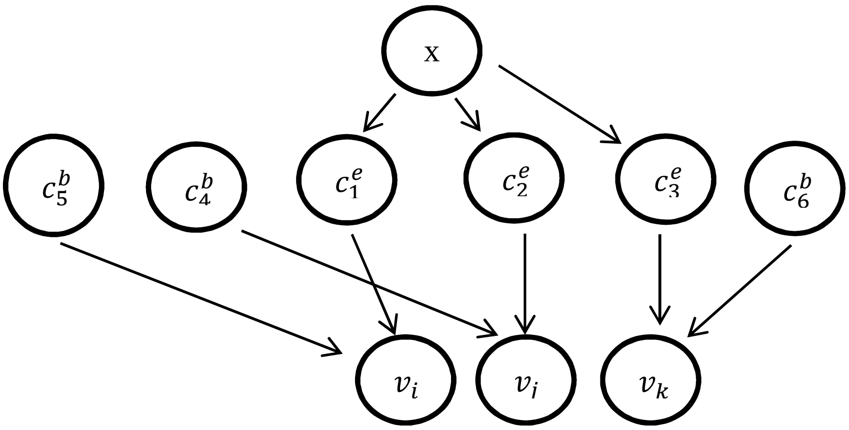

The posterior competence of the three agents is as follows:

$$p_i^1 = \Pr(v_i = x | x, c_1^b, c_5^e) = 0.6$$

$$p_j^1 = \Pr(v_j = x | x, c_2^b, c_4^e) = 0.6$$

$$p_k^1 = \Pr(v_k = x | x, c_3^b, c_6^e) = 0.6$$

In other words, the combined effect of the two causal factors, received privately by each agent, is to increase the competence level of each agent by +0.1. If the social choice is decided by the aggregation procedure of majority rule then the probability of a correct majority verdict given the votes of i, j, k is $P = 0.648$.

Case 2

Now consider a slightly different situation as represented in figure 7.3 below.

Figure 7.3: a causal network with one common package of information.

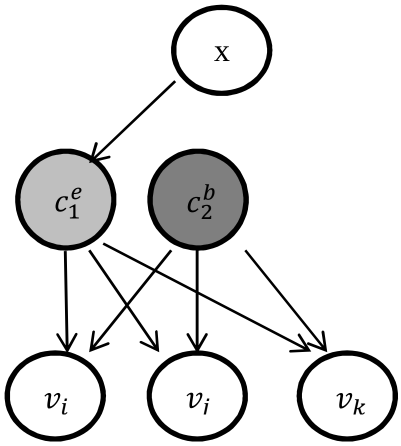

Here all the agents receive the same two pieces of information. The prior competence of agent i will be $p_i^0 = \Pr(v_i = x|x) = 0.5$. We need to update the competence of agent i in light of the information that they have received. The posterior competence of agent i is $p_i^1 = \Pr(v_i = x|x, c_1^e, c_2^b) = 0.6$. Agents j and k are not independent of agent i since they have common causal factors. Once we know the way agent i votes this will increase the probability that agents j and k vote in the same way. Formally, $p_j^1 = \Pr(v_j = x|x, v_i) > p_j^o = \Pr(v_j = x|x)$. However once we conditionalise on the common factors, we regain independence between the three agents. The posterior competencies of agents j and k conditional on the evidential and background information they receive are as follows:

$$p_j^1 = \Pr(v_j = x|x, c_1^e, c_2^b) = 0.6$$

$$p_k^1 = \Pr(v_k = x|x, c_1^e, c_2^b) = 0.6$$

We can see that independence has been regained by conditionalising on common factors by noting, inter alia, that $p_j^2 = \Pr\left(v_j = x \middle| x, c_1^e, c_2^b, v_i\right) = p_j^1 = \Pr\left(v_j = x \middle| x, c_1^e, c_2^b\right)$=0.6. [116]

Dietrich and Spiekermann (unpublished b) provide a general form of Reichenbach's common cause principle:

"**Common Cause Principle** (stated informally). Phenomena which do not causally affect each other are probabilistically independent conditional on their common causes." (p.5)

In the example above the votes of agents do not causally affect each other. We have conditionalised on the common causes of agent's votes, so the votes of agents will be probabilistically independent.

If the social choice in case 2 immediately above is decided by the aggregation procedure of majority rule then the probability of a correct majority verdict given the votes of i, j, k is $P = 0.648$.

In case 1 there were three packages of background and evidential information which each had a truth-conducive value of +0.1 (that each increase an agent's competence level by a value of 0.1). In the case 2 there was one package of background and evidential information that had a truth-conducive value of +0.1. There is three-times as much information available to the group in case 1 as

[116] And if agent i votes for alternative $\neg x$ it makes it neither more nor less likely that agent j will vote for alternative $\neg x$.

there is in case 2 and yet the probabilities of a majority verdict in each case are identical. It seems as if the distribution of information among agents in case 1 is suboptimal.

Case 3

Suppose we face the same set of causes as in case 1. However in case 3, instead of dividing up the information among different agents, all information is common. We can represent this in the figure below:

Figure 7.4: a causal network with three common packages of information.

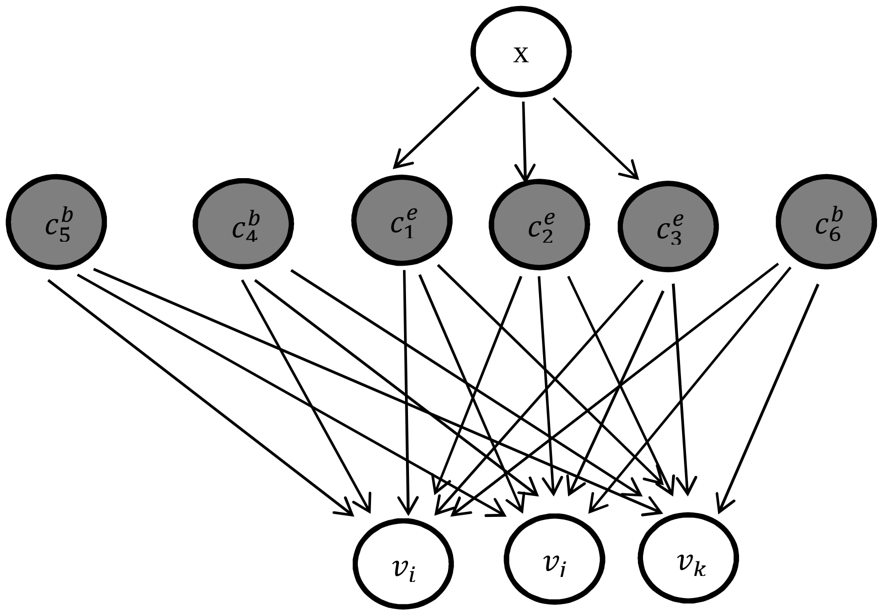

The posterior competence of the three agents is:

$$p_i^1 = \Pr(v_i = x | x, c_1^e, c_2^e, c_3^e, c_4^b, c_5^b, c_6^b) = 0.77$$

$$p_j^1 = \Pr(v_j = x | x, c_1^e, c_2^e, c_3^e, c_4^b, c_5^b, c_6^b) = 0.77$$

$$p_k^1 = \Pr(v_k = x | x, c_1^e, c_2^e, c_3^e, c_4^b, c_5^b, c_6^b) = 0.77$$

Because we have conditionalised on all the common factors, the votes of agents i, j and k are independent. As can be seen, there is randomness in each of their votes. This randomness comes from whatever competence-generating factors the agents hold privately. If the social choice is decided by the aggregation procedure of majority rule then the probability of a correct majority verdict given the votes of i, j, k is $P = 0.86563$. There are clear epistemic advantages to sharing truth-conducive information.

In an earlier section on premise-based procedures I argued that the increased epistemic performance generated by taking a premise-based approach is due in part to the competence of agents being modular across propositions. The modularity of competence across propositions means the vote of a single agent is independent across propositions. The same truth-conducive background factor, such as experience of paying close attention to complicated testimony, may generate high levels of competence for an agent in several different social choice problems.

The examples of information sharing, immediately above, show that competence can also be modular across agents. The same competence-generating factors of background and evidential information can be shared by different agents and in the process boost their levels of competence, while preserving independence conditional on the common factors.

We can explain the modularity of competence across agents by analogy with the coin-toss. Tossing the same flawed coin with a 0.51 bias 100 times is equivalent to tossing 100 identical coins with a 0.51 bias produced from the same flawed die at the mint[117]. The flaw in the die at the mint which places a +0.01 bias towards heads on each coin is analogous to the truth-conducive background factor of experience of paying close attention to complicated testimony, which places a +0.1 bias towards the truth on the judgements of agents. Agents can share this same truth-conducive background factor and yet still be independent.

There is one final point to be said in favour of the proposal of sharing information among agents. The classic CJT requires that the competence levels of agents are homogeneous and that agents are independent conditional on the state of the world. Neither of these assumptions is plausible in real-world social choice problems. When agents share information we exchange independence conditional on the state of the world with independence conditional on the state of the world and common factors[118]. If agents share information then their levels of competence conditional on the common factors will be homogeneous, as per the classic CJT. The sharing of information between agents provides a justification for the assumption of homogeneous levels of competence.

[117] As Estlund (2008) states "Now obviously the same would be true if instead of one coin flipped repeatedly, we consider many coins, all weighted the same way, each having a 51 percent chance of coming up heads." (p.224)
[118] See Dietrich and Spiekermann (forthcoming a,b)

Problems with information sharing

I will present two concerns with the solution to the problem of finite information of information sharing: it may not be possible to share all the truth-conducive causal factors and there remains a possibility of misleading information.

I have argued that the reason there is still some independence in the votes of agents, conditional on common factors, is that at least some of the competence-generating factors are held privately and not shared with other agents. There may be many types of truth-conducive competence-generating factors which, from the perspective of group epistemic performance it is desirable that agents share, but which they are unable to share. For example, while background factors such as textbooks on pathology or courses in how to be a juror, can be shared between agents it may be that the background factor of experience of paying close attention to complicated testimony is just the sort of background competence-generating factor that *cannot* be shared between agents. This background factor is built up over a lifetime and cannot be passed on to other jurors during the trial. Similarly, while evidential factors such as fingerprint evidence may be shared between agents, a witness may not be able to share with jurors the evidential factor of the precise colour of the shirt the defendant was wearing.

Secondly, there remains an intuitive concern with the argument that competence-generating factors can be shared by agents and boost the competence levels of agents, while retaining some independence in the agent's votes. I have argued that if we have a group with competence levels $p_i, p_j, p_k = (0.6,0.6,0.6)$ then it makes no difference to the probability of a correct majority verdict whether their competence levels were generated by three private packages of information (case 1) or whether they share the same package of information (case 2). But surely there must be *some* epistemic advantage to having three different sources of information. We can account for this intuition in part, as I have done above, with the assertion that there is just more information in the group with three different packages of information than there is in the group with one package of information. If the three different packages are held privately then this is just a sub-optimal distribution of information.

There is however a more significant concern with the competence of agents being generated by the same truth-conducive factors. Thus far I have assumed that the competence-generating factors are instantiations of random variables. If, for example, an agent has the background factor of experience of paying close attention to complicated testimony then it was assumed that this had a truth-conducive influence on an agent's level of competence. However we can also treat causal factors as random variables that can take two values. For example, the factor of experience of paying close attention to complicated testimony could either be truth-conducive and increase an agent's competence; or the factor of experience of paying close attention to complicated testimony could be misleading in which case it *decreases* an agent's level of competence. If our group of three agents shares the same package of information (case 2) and this turns out to be misleading then each agent will have competence $p_i, p_j, p_k < 0.5$. If however each of our three agents has a different, private package of information (case 1) then it will only be the case that the competence of each agent is $p_i, p_j, p_k < 0.5$ if each of these packages of information turns out to be misleading. It is more likely that a single package of information turns out to be misleading than for three separate packages of information to turn out to be misleading. Separate sources of evidence hedge against the possibility of misleading evidence.

We can use sample calculations to consider the differences between cases 1, 2 and 3 when we allow for the packages of information to be truth-conducive (increase competence) or misleading (decrease competence).

Case 1: three private packages of information

- If all three packages of information are truth-conducive then the competence of all agents is $p_{i,j,k} = 0.6$ and the probability of a correct majority verdict is $P = 0.648$.
- If one of the packages of information is misleading then the competence of agents is $p_i, p_j, p_k = (0.4, 0.6, 0.6)$ and the probability of a correct majority verdict is $P = 0.552$. There are three ways in which this situation could arise: the information of the first, second or third agent could be misleading.
- If two of the packages of information are misleading then the competence of agents is $p_i, p_j, p_k = (0.4, 0.4, 0.6)$ and the probability of a correct majority verdict is $P = 0.448$. There are three ways in which this situation could arise: the information of the first, second or third agent could be truth-conducive.

• If all three packages of information are misleading then the competence of all agents is $p_i, p_j, p_k = 0.4$ and the probability of a correct majority verdict is $P = 0.352$

Let C_1 represent the case where the first package of information is truth-conducive and $\neg C_1$ represent the case where the first package of information is misleading. $MC|C_1$ represents the case that there is a majority for the correct alternative given that the first package of information is truth-conducive. The probability of a correct social choice is given by:

$$[\Pr(C_1\&C_2\&C_3) \times \Pr(MC \mid C_1\&C_2\&C_3)] + 3[\Pr(C_1\&C_2\&\neg C_3) \times \Pr(MC \mid C_1\&C_2\&\neg C_3)] + 3[\Pr(C_1\&\neg C_2\&\neg C_3) \times \Pr(MC \mid C_1\&\neg C_2\&\neg C_3)] + [\Pr(\neg C_1\&\neg C_2\&\neg C_3) \times \Pr(MC \mid \neg C_1\&\neg C_2\&\neg C_3)]$$

Suppose the probability of a given package of information being misleading is 0.6. The probability that a group with three private packages of information choose the correct alternative via majority rule is:

$$(0.6^3 \times 0.648) + (3(0.6^2 \times 0.4) \times 0.552) + (3(0.4^2 \times 0.6) \times 0.448) + (0.4^3 \times 0.352) = 0.52998$$

Case 2: one common package of information

- If the information is truth-conducive then the competence of each agent will be $p_i, p_j, p_k = 0.6$ and the probability of a correct majority verdict will be $P = 0.648$.
- If the information is misleading then the competence of each agent will be $p_i, p_j, p_k = 0.4$ and the probability of a correct majority verdict will be $P = 0.352$

Suppose the probability of a given package of information being misleading is 0.6. The probability that a group with one common package of information choose the correct alternative via majority rule is:

$$(0.6 \times 0.648) + (0.4 \times 0.352) = 0.5296$$

Case 3: three common packages of information

- If all three packages of information are truth-conducive then the competence of each agent will be $p_i, p_j, p_k = 0.77$ and the probability of a correct majority verdict will be $P = 0.86563$
- If two of the packages of information are truth-conducive and one of the packages of information is misleading then the competence of each agent will be $p_i, p_j, p_k = 0.6$ and the probability of a correct majority verdict will be $P = 0.648$. There are three ways in which this situation could occur: the first package of information is misleading, or the second is misleading, or the third is misleading.

- If one of the packages of information is truth-conducive and two of the packages of information are misleading then the competence of each agent will be $p_i, p_j, p_k = 0.4$ and the probability of a correct majority verdict will be $P = 0.352$. There are three ways in which this situation could occur: the first, or second, or third package of information is truth-conducive.
- If all three packages of information are misleading then the competence of all agents is $p_i, p_j, p_k = 0.23$ and the probability of a correct majority verdict is $P = 0.13437$

Suppose the probability of a given package of information being misleading is 0.6. The probability that a group with three common packages of information choose the correct alternative via majority rule is:

$$(0.6^3 \times 0.896) + (3(0.6^2 \times 0.4) \times 0.648) + (3(0.4^2 \times 0.6) \times 0.352) + (0.4^3 \times 0.104) = 0.5815$$

Figure 7.5: the probability of a correct majority verdict.

Probability that a package of info. is truth-conducive	**Case 1: three private packages**	**Case 2: one common package**	**Case 3: three common packages**
0	0.352	0.352	0.13437
0.1	0.38102	0.3816	0.20185
0.2	0.41043	0.4112	0.2731
0.3	0.44013	0.4408	0.34716
0.4	0.47002	0.4704	0.42311
0.5	0.5	0.5	0.5
0.6	0.52998	0.5296	0.57689
0.7	0.55987	0.5592	0.65284
0.8	0.58957	0.5888	0.7269
0.9	0.61898	0.6184	0.79815
1.0	0.648	0.648	0.86563

What can we infer from these sample calculations? There are two questions that we should ask. Firstly, given that each agent receives a fixed amount of information, is it best if this information is private or held in common with other agents? This first question is addressed by comparing the results for case 1 (each agent receives a fixed, private amount of information) and case 2 (each agent receives a fixed, common amount of information). The second question to ask is given a fixed amount of information, is it best if this information is held in common between agents or parcelled out in private packages? This second question is addressed by comparing case 1 (the information is parcelled out in private packages) and case 3 (the information is held in common between agents).

The results for case 1 (where each of three agents receives one private package of information) are *very* similar to case 2 (where each of three agents shares one common package of information). If anything, case 1 is slightly *less* reliable when information tends to be misleading but slightly *more* reliable when information tends to be truth-conducive. Intuitively there are two forces that pull in opposite directions. If agents have different packages of information generating their competence levels we hedge against the possibility of misleading information. If a piece of information does turn out to be misleading it will only affect the competence of a single agent. On the other hand, where agents have different factors generating their competence there is just more information in the group and so there is a greater chance of some of that information being misleading.

The comparison between case 3 and case 1 is much clearer. If the packages of information tend to be truth-conducive then it is best if shared among agents (case 3); if the packages of information tend to be misleading then it is best if they are held privately (case 1). Again, if packages of information tend to be misleading then by parcelling them out to different agents we quarantine their impact on agent competence levels. If packages of information tend to be truth-conducive then the impact of this information is maximised by applying it multiple times to different agents.

Summary

Given finite amounts of information, what is the best institutional response to maximise the probability of a correct social choice? The choice of aggregation procedure can make a significant difference. In particular, weighted majority rule or expert dictatorship will maximise the probability of a correct social choice. If it is possible to disaggregate a social choice into premises then employing a premise-based procedure with majority rule (or weighted majority rule) can also make a significant difference.

It is the sharing of information between agents that can have the greatest impact on a group's ability to identify the true state of the world. When the information in a group is truth-conducive there are obvious advantages to sharing this information between agents and there are no disadvantages to sharing information. If there is a possibility that the information might be misleading, then provided the information is more likely to be truth-conducive than misleading, again the information should be shared.

The mechanism that accounts for the boost to group epistemic performance that comes from a premise-based procedure or the sharing of information is that the competence of agents is a modular epistemic capacity. As such it is possible to "reapply" the truth-conducive impact of finite amounts of information.

Suppose that as the result of a search procedure a group of three agents has identified truth-conducive information of a strength such that the competence of agents increases from $p_i^0 = 0.5$ to $p_i^1 = 0.6$. If post-search this information is shared by the three agents the competence levels of each of these agents will be $p_i^2 = 0.77$. If the social choice is determined by the aggregation procedure of majority rule then the probability of a correct social choice will be:

$$P_{Majority\ Rule} = 0.86563$$

If this same information is shared among 11 agents then the probability of a correct majority verdict will be:

$$P_{Majority\ Rule} = 0.977$$

If this same information is shared among 101 agents then the probability of a correct majority verdict will be[119]:

$$P_{Majority\ Rule} = 1.0$$

The fact that each agent has some private, background, competence-generating factors means that votes of the agents will be independent. And the fact that these private, background, competence-generating factors are unknown to the social planner means the competence levels of agents are less than 1.0, that there is still some randomness in the votes of agents.

Provided that each agent in the group has some private background factors and provided that the agents share the known truth-conducive evidential and background factors, then there only needs to be a small amount of truth-conducive information (evidential and background) for a group to come extremely close to identifying the true state of the world using majority rule. As we have seen in the sample calculation immediately above, all we need is evidence of +0.27 truth-conducive value and 101 agents with some private

[119] Allowing for rounding.

background competence-generating factors for the group to be close to certain to identify the true state of the world.

Chapter 8: Conclusion.

This thesis has addressed the social epistemic mechanisms operating in groups of political agents; the institutional arrangements employed by groups of political agents that allow the group to track the truth. I have argued that social choices in political settings occur via a two-staged process. Firstly, there are search procedures by which agents find truth-conducive information. Secondly, there are aggregation procedures by which agents pool the information they have found. I have identified five social epistemic mechanisms that can operate during this two-staged process: two that can operate during the search procedure and three that can operate during the aggregation procedure. For each of these social epistemic mechanisms, increasing group size is epistemically virtuous.

During the search procedure there are, firstly, the institutional arrangements as captured by the Spatial Search Theorem. If every agent has at least some possibility of moving to the location of a piece of information and each agent searches some different locations then as the size of the group increases, the probability of finding the piece of information also increases. Relatedly, as the size of the group increases the amount of information identified by the group increases. In the introduction I highlighted the taxonomy of models of potential group productivity presented in Steiner (1966). This first social epistemic mechanism can be characterised as an additive model of group productivity. As the group size increases the probability the object is found is the sum of the probabilities that individual agents in the group find the object.

The second social epistemic mechanism operating during the search procedure is captured by the Search Recognition Theorem. If each agent has some possibility of recognising an object at a particular location and the recognition capacities are independent then as the number of agents visiting the location increases the probability of at least one agent recognising the object also increases. This second social epistemic mechanism can again be characterised as an additive model of group productivity, according to the Steiner taxonomy.

In this thesis I considered three types of aggregation procedures, the institutional features of which amount to social epistemic mechanisms. Firstly, expert dictatorship can be characterised as a disjunctive model of group epistemic productivity. Under dictatorship the probability the group identifies the true state of the world is limited to the competence of the most competent member of the group. If competence in the wider population is heterogeneous then as group size increases the competence of the most competent member of a group should increase. Unanimity rule can be characterised as a conjunctive model of group epistemic productivity. Increasing group size generally decreases the probability of a unanimous choice, since all agents must perform the same action and the probability of them all doing so decreases with group size. However negative reliability unanimity rule is more and more likely to track the truth as group size increases since the probability of all agents voting for the incorrect alternative decreases as group size increases. Negative reliability unanimity rule can be characterised as a disjunctive model of group productivity since it only requires a single agent to vote for the correct alternative for the group to avoid the incorrect alternative as the social choice.

The institutional features of the aggregation procedure of majority rule mean a group using majority rule are increasingly likely to track the truth as group size increases. Therefore, the institutional features of majority rule also amount to a social epistemic mechanism. Majority rule can be classed as a compensatory model of group productivity, according to the Steiner taxonomy. Although some agents may make a mistake and vote for the wrong alternative, their votes can be offset by other agents voting for the correct alternative. If agents are sufficiently competent then as group size increases it becomes increasingly likely that there will only be a minority voting for the wrong alternative, whose votes are offset by a majority voting for the correct alternative.

The Steiner taxonomy of models of group productivity cannot be thought of as being exhaustive. For example, there may be further models of group productivity that are exponential in nature- as the number of agents increases the institutional features of an aggregation procedure may mean that the probability of a correct social choice increases rapidly. Similarly the five different social epistemic mechanisms I have identified as operating during the two stages of search then aggregation should not be thought of as exhaustive. There are 2^{2^n} possible aggregation procedures for a dichotomous choice[120] and conceivably many more than three of these possible aggregation procedures will have institutional features that allow a group to track the truth as group size increases. There may also be other, different, models of search procedure whose institutional features allow a group to find the objects of search.

[120] Christian List, unpublished lecture notes.

Furthermore there may be social epistemic mechanisms, formal or informal institutional arrangements in groups of agents, that allow the groups to track the truth, that do not fit within the two-staged framework of search and aggregation. For example, Vermeule (2009) proposes, inter alia, an evolutionary account of 'many-minds' arguments in legal theory. Here, the judgements of a group of agents will, over time, weed out unfit policies through a process of evolution. In a political setting, the combined wisdom of a series of parliaments should, over time, work to weed out unsatisfactory parts of the law. This iterative improvement does not fit into the one-shot framework of search followed by aggregation that I propose. Similarly, some types of deliberation may be successful at tracking the truth as group size increases, but deliberation is not a necessary component of the two-staged framework of search and aggregation procedures.

Interestingly, the institutional arrangements in the five social epistemic mechanisms identified seem to be of two kinds: in some cases we want to encourage agents to act differently and to capture the influence of outliers. In other cases we want the agents to act in a similar fashion so as to weed out the influence of outliers. A spatial search procedure requires agents to perform different tasks and relies on an agent visiting a location no other agent visits. Increasing the size of the group increases the probability the group will contain such an exceptional agent. The search recognition procedure requires agents to perform a similar task. As we increase the size of the group the probability that one of the group performs the task properly (recognises the object) increases.

The five social epistemic mechanisms can provide epistemic justifications for various democratic virtues, including widening participation in political decision making, free speech and freedom of association, and for diversity or pluralism in the population. By increasing the size of the group participating in political decision making we increase the probability of finding particular pieces of evidential and relevant background information. We also increase the total amount of information available to the group. Freedom of association prevents arbitrary restrictions on the search procedures of individual agents. Free speech can alert fellow agents to locations containing new or difficult to recognise pieces of information. In general it is best to include as many varied agents as possible as the variety of initial partitions, locational conventions, start points and search heuristics means these group members are able to identify different pieces of information.

Increasing the size of a group participating in political decision making is also epistemically virtuous when it comes to the aggregation of judgements, no matter which of the three aggregation procedures are used. Increasing group size increases the probability of a correct majority verdict, a correct judgement of an expert dictator and increases the probability of avoiding a unanimous verdict for the incorrect social choice (given appropriate competence levels and independence relations).

The only model of potential group productivity (in this case epistemic productivity) identified by Steiner and not occurring in the two-staged framework of search and then aggregation procedures is a complementary model. Under a complementary model different parts of a task are performed by different agents. In an epistemic setting, different propositions whose conjunction deductively entails a conclusion may be possessed by different agents. It was argued in chapter 6 of the thesis that the post-search distribution of information across agents could mean different agents possess partial pieces of information. The partial pieces of information may be sufficient for the group to determine the true state of the world. However none of the aggregation procedures I have presented can appropriately pool these pieces of information, since the partial pieces of information do not directly increase the competence levels of agents. In such cases neither dictatorship, nor negative reliability unanimity rule, nor majority rule will correctly identify the true state of the world. The information dispersed among agents can only be pooled into a correct social choice by sharing it directly, perhaps via deliberation.

I also argued in chapter 7 that the truth-conducive impact of finite information is maximised by sharing that information between agents. A period of deliberation, post search but prior to aggregation, would allow for the sharing of this information.

Unfortunately there are a number of potential problems with deliberation and information sharing, including information cascades, group think and group polarisation. I will touch on these very briefly.

Sunstein (2002) discusses the phenomena of 'group polarisation' where the judgements of agents post-deliberation are more extreme than their pre-deliberation judgements. A juror who pre-deliberation supports a fairly long prison sentence for an offender can support an extremely long sentence after discussing the issue with their fellow jurors. Sunstein cites three possible explanations for group polarisation: an initial skew in the makeup of a group, a desire to fit in with the group and overconfidence of agents with extreme views.

Closely related to group polarisation is ‘group think’, as studied by Janis (1972), where the group engaging in deliberation excludes information that would disrupt the consensus. Group think can lead to a group making the wrong decision even if the judgements of the group are not subjected to the shift seen in group polarisation. Group think could be caused by the homogeneity of the group and the desire for cohesion. An agent might withhold from the group a proposition that could disrupt a consensus on the (incorrect) alternative. Agents may not contribute their partial pieces of information identified during their searches if doing so would disrupt the consensus.

Finally, information cascades can occur when agents ignore their own private information and instead base their judgements on the judgements of other agents, who in turn based their judgements on the judgements of other agents. Information cascades can be problematic since although it seems as if an agent’s judgement is based on the large body of information contained in the judgements of previous agents, the amount of information may in fact be quite small (see Goodin and Spiekermann, 2011).

Identifying the social epistemic mechanisms operating in deliberative practices (the institutional arrangements for discussion and debate that facilitate information sharing) while avoiding the traps of information cascades, group think and group polarisation, is the focus of the next phase of my research.

Bibliography.

Anderson, E. (2006) 'The Epistemology of Democracy', *Episteme*, 3: 9-23.

Ben-Yashar, R. and Paroush, J. (2000) 'A Non-asymptotic Condorcet Jury Theorem', *Social Choice and Welfare*, 17: 189-199.

Bovens, L. and Rabinowicz, W. (2006). 'Democratic Answers to Complex Questions – an Epistemic Perspective' *Synthese* 150 (1):131-153.

Bradley, R. (2006) 'Taking Advantage of Difference of Opinion', *Episteme*, 3(3): 141-155.

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

Cohen, J. (1986) 'An Epistemic Conception of Democracy', *Ethics*, 97(1): 26-38.

Copp, D. (1999) 'Can political truth be a hazard for democracy?' in Copp, D., Hampton, J. and Roemer, J.E. *The Idea of Democracy*, Cambridge University Press.

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

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

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

Edelman, P.H. (2002) 'On Legal Interpretations of the Condorcet Jury Theorem', *Journal of Legal Studies*, 31:327-349.

Estlund, D.M., Waldron, J., Grofman, B. and Feld, S.L. (1989) 'Democratic Authority and the Public Interest: Condorcet and Rousseau Revisited', *American Political Science Review*, 83(4): 1317-1340.

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

Estlund, D.M. (1997) 'Beyond fairness and deliberation: the Epistemic Dimension of Democratic Authority', in Bohman, J. and Rehg, W. (eds.) *Deliberative Democracy*, MIT Press.

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

Fuerstein, M. (2008) ‘Epistemic Democracy and the Social Character of Knowledge’, *Episteme*, 5(1): 74-93.

Goldman, A.I. (1999) *Knowledge in a Social World*, Oxford University Press.

Goldman, A.I. (2010) ‘Systems-Oriented Social Epistemology’, in Gendler, T.S. and Hawthorne, J. *Oxford Studies in Epistemology: Volume 3*, Oxford University Press.

Goodin, R.E. and Spiekermann, K. (2011), ‘Courts of Many Minds’, *British Journal of Political Science*, FirstView Article.

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.

Janis, I. (1972). *Victims of groupthink; a psychological study of foreign-policy decisions and fiascoes*. Boston: Houghton.

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

Kaniovski, S. (2010) ‘Aggregation of Correlated Votes and Condorcet’s Jury Theorem’, *Thoery and Decision*, 69(3): 453-468.

Kitcher, P. (1990) ‘The Division of Cognitive Labour’, *Journal of Philosophy*, 87(1): 5-22.

Kitcher, P. (1993) *The Advancement of Science: Science Without Legend, Objectivity Without Illusions,* Oxford University Press.

List, C. (2006) ‘The discursive dilemma and public reason’, *Ethics*, 116(2): 062-402.

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

List,C. and Goodin, R.E. (2006) ‘A Conditional Defence of Plurality Rule: Generalizing May’s Theorem in a Restricted Informational Environment’, *American Journal of Political Science*, 50(4): 940-949.

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.

List, C. and Vermeule, A. (2010) ‘Independence and Interdependence: Lessons from the Hive’, Harvard Public Law Working Paper No. 10-44

Mill, J. S. (1861) *Considerations on Representative Government*, online.

Muldoon, R. and Weisberg, M. (2011) ‘Robustness and Idealization in Models of Cognitive Labour’, *Synthese*, 183(2): 161-174.

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.

Peter, F. (2008) ‘Pure Epistemic Proceduralism’, *Episteme*, 5(1): 33-55.

Plato (1988) *Republic*, R Waterfield translation, OUP.

Romeijn, J-W and Atkinson, D. (2011) ‘Learning juror competence: a generalised Condorcet Jury Theorem’, *Politics, Philosophy and Economics,* 10(3): 237-262.

Rousseau, J.J. (1762) *The Social Contract*, online.

Sunstein, C.R., (2002) ‘On a Danger of Deliberative Democracy’, *Daedlus*, 131(4): 120-124.

Steiner, I.D. (1966) ‘Models for Inferring Relationships between Group Size and Potential Group Productivity’, *Behavioural Science*, 11(4): 273-283.

Vermeule, A. (2009) ‘Many-Minds Arguments in Legal Theory’, *Journal of Legal Analysis*, 1(1): 1-45.

Young, H.P. (1988) ‘Condorcet’s Theory of Voting’, *American Political Science Review*, 82(4): 1231-1244.

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

Wittgenstein, L. (1953) *Philosophical Investigations*, Blackwell Publishing.

Wolff, J. (1994) 'Democratic Voting and The Mixed Motivation Problem', *Analysis*, 54(4) 193-196.

Wolff, J. (1996) *An Introduction to Political Philosophy*, Opus.